

eBook
April 2023

End-to-End High-Speed Digital Design

SPONSORED BY



Table of Contents

3

Introduction

Patrick Hindle
Media Director, Signal Integrity Journal/Microwave Journal

5

How To Choose an Interconnect for PCIe 6.0 High-Speed Systems

Steve Krooswyk , Anthony Fellbaum, and Matt Burn

8

Compute Express Link (CXL) A New Coherency Protocol Over PCIe

Keysight

17

USB4 Version 2.0 from Simulation to Tx, Rx and Interconnect Test

Keysight

21

USB-C Signal Analysis Methodology

Brian Holden
Kandou Bus

24

DDR5 Signal Integrity Fundamentals

Tim Wang Lee
Keysight

27

Back to Basics: IBIS/IBIS-AMI and the Path to LPDDR5

HeeSoo Lee and Fangyi Rao
Keysight

33

Managing PCB Crosstalk

Donald Telian
SI Consultant, SiGuys

37

Proper Ground Return Via Placement for 40+ Gbps Signaling

Michael Steinberger, Donald Telian, Tsuk Michael, Iyer Vishwanath, and Yanamadala Janakinadh
MathWorks and SiGuy

Introduction

End-to-End High-Speed Digital Design

Technology is advancing faster than ever, creating new opportunities and challenges for PCB design. This acceleration results from three main factors: faster data transfer and processing for cloud computing, AI, gaming, and VR; smaller, lower-power, and denser chips from improved semiconductor fabrication; and new standards/protocols for higher device and platform compatibility.

High-speed digital links at 224 Gbps per lane, device miniaturization, and ultra-low power budgets result in sophisticated systems that are more complex to optimize signal integrity performance and reliability. Engineers require close collaboration and seamless flows from concept to simulation, emulation, and testing to keep up with technology's pace and meet shrinking time-to-market expectations.

This eBook will cover how to connect design workflows to make end-to-end high-speed digital design a reality. It will discuss the challenges and trade-offs associated with optimizing design performance for technologies such as PCIe 6.0, USB4 Version 2.0, and USB-C, DDR5, and LPDDR.

From silicon development and interconnect validation to low-level PCB layout for vias and signal lines, the holistic methodology with an integrated software platform democratizes electronic design automation while streamlining time-to-insight.

The eBook starts with PCIe 6, discussing interconnect design choices and a must-read overview of the Compute Express Link (CXL) technology, which introduces new high bandwidth and low latency protocols on top of PCIe for attaching accelerators and memory buffers.

Next, we dive into USB4v2 with the foundation blocks for an end-to-end digital design approach and best-in-class signal analysis methodology to meet strict interoperability requirements in the vast USB Type-C ecosystem.

Then we switch gears to DDR5 signal integrity fundamentals and how IBIS/IBIS-AMI is the fastest and easiest way to validate designs avoiding multiple boards spins as we forge the path towards LPDDR5.

Lastly, the book closes out with best practices for managing PCB crosstalk and proper placement of ground return as we move beyond 40+ Gbps signaling.

We are confident that this eBook will provide new insights addressing the challenges of next-gen high-speed digital design and the benefits of implementing the end-to-end approach that connects workflows from concept to emulation, simulation, and test.

Patrick Hindle, Media Director, Signal Integrity Journal/Microwave Journal



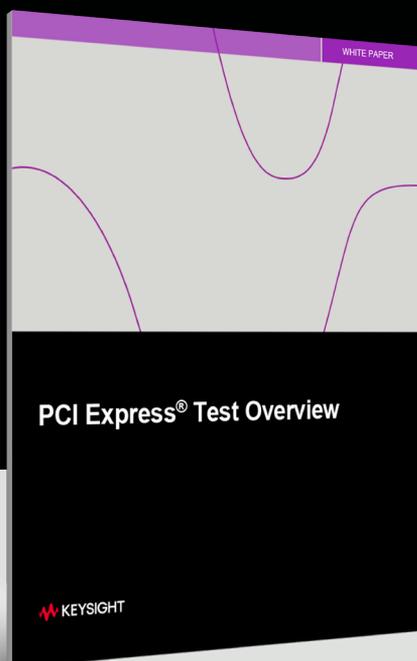
Unlock Success in Your Transition to PCIe® 6.0

Read Keysight's newest white paper to quickly learn how to overcome your toughest test challenges with PCIe® 4.0, 5.0, and 6.0. Increase your knowledge of complete testing solutions:

- system simulation
- interconnect design
- transmitter test
- receiver and link equalization test
- protocol layer test

Explore how to test the future of PCIe technology with Keysight's continuous innovations.

Get the white paper



PCI-SIG®, PCIe® and the PCI Express® are US registered trademarks and/or service marks of PCI-SIG.

Innovators
▶ start here

Learn more at www.Keysight.com

How To Choose an Interconnect for PCIe 6.0 High-Speed Systems

Steve Krooswyk , Anthony Fellbaum, and Matt Burns

Simply put: We need more bandwidth. In this era of extreme data collection and processing, we're now designing systems using SERDES at 112 Gbps PAM4 (28 GHz Nyquist rate) per lane and looking at 224 Gbps and beyond. One of the key challenges in implementing these high data rates is choosing the interconnect that will maintain appropriate signal integrity, preferably with some margin. In this article, we'll go over some interconnect design choices in the context of the new PCI Express® (PCIe®) 6.0 specification and provide suggestions on how to overcome design tradeoffs and challenges.

FINDING A STANDARD

In the high-tech industry, the Peripheral Component Interconnect Special Interest Group (PCI-SIG®) was founded in 1992. For more than 30 years, PCI-SIG member companies have championed PCI® and PCIe technology as the de facto interconnect in any number of high-performance computing applications.

The sixth generation of PCIe technology, fully released in January 2022, again doubles the bandwidth (64 GT/s) compared to the previous generation. Additionally, PCIe 6.0 offers processor-agnostic, cost-effective, power-efficient, low-latency, scalable connectivity between components while maintaining backward compatibility. These key features help PCIe 6.0 support the insatiable bandwidth demand in artificial intelligence, machine learning, networking, communication systems, storage, crypto mining, high-performance computing applications, and more. In addition, it is driving the need for innovative interconnect systems to handle the increased data rates.

NEW SPEC: NEW CHALLENGES

Doubling the data rate for PCIe 6.0 presents challenges for two-level signaling, where a suggested 32 GHz operating frequency and associated insertion loss and noise reductions put an extreme burden and expense on PCB design and manufacturing. For the last few decades, high-speed links in the data center and compute environment have used two voltage levels: low voltage representing binary 0 and high for 1 (non-return to Zero (NRZ) or PAM2). As frequency increased to achieve more bandwidth, loss, and reflection became a challenge. The solution, four-level pulse amplitude modulation (PAM4), enables PCIe 6.0 to double the data without increasing the operating frequency.

Operating at a decreased signal amplitude for each bit (1/3rd), the PAM4 link experiences a 9.54 dB reduction in signal-to-noise ratio. To achieve the reduction in noise, some improvements were necessary. For example, PCIe 6.0 includes advances in receiver equalization to compensate for frequency-dependent loss, with an increase in decision feedback equalization (DFE) from 3-taps to 16-taps. In addition, loss budget requirements were reduced for larger root complex packages, and the channel target was reduced from 36 to 32 dB.

NEXT GENERATION CONNECTORS

Similarly, we expect loss and noise requirements for PCIe 6.0 CEM connectors to tighten. Although the 6.0 CEM specification is not yet released, we can expect new metrics like integrated crosstalk noise (ICN) to maintain an important presence in PCIe 6.0 CEM specification. ICN is a weighted, single-number metric used to test the overall connector crosstalk for compliance.

To perform this test, a weighting function (shown on left in Figure 1) is defined that is proportional to the power transferred through a PCIe 32 GT/s NRZ or 64 GT/s PAM4 system and is then used to filter the measured connector crosstalk power sum. (Resonances and other minor excursions beyond the limit line that have little relevance to actual system bit-error-rates will be permitted.)

As the insertion loss target has been reduced in the PCIe 6.0 specification, there is also a need to improve PCB material to meet sufficient channel reach, targeting near or below 1.0 dB/in at 16 GHz.

CABLE SOLUTIONS RISE TO THE CHALLENGE

To satisfy the more stringent loss targets in the PCIe 6.0 specification, designers can consider a shift from PCB transmission line to cabled solutions. In addition to improving loss performance, this approach can also extend reach. For instance, in **Figure 1**, the Samtec PECFF Emulation Platform illustrates configurable cable topologies between mock GPU cards on a SFF-TA-1002 based backplane in a typical artificial intelligence (AI)/machine learning (ML) system architecture. In this case, the trade-off ratio is 10:55 in. (PCB:Cable) using 34 AWG Eye Speed® Twinax in the chassis.

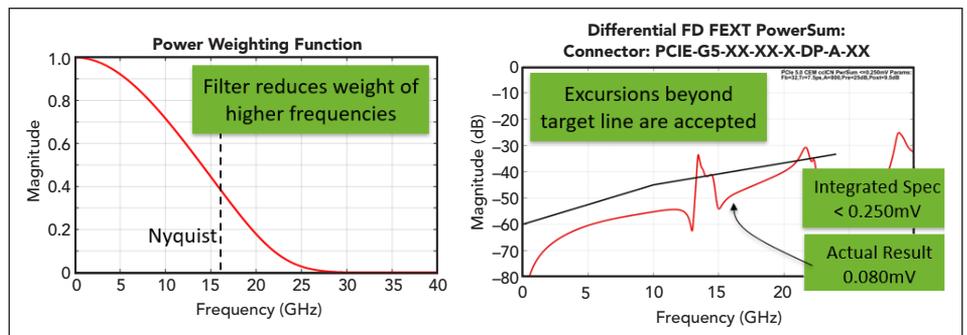
LOSS BUDGETS

System architecture decisions (PCB and cable length, whether a repeater is needed, etc.) can be made without extensive simulation expertise by using a loss budget with appropriate estimates. However, raw arithmetic assumptions of system losses do not consider the complexity of increased noise sensitivity of the PAM4 link, so they can lead to optimistic and potentially disastrous assumptions. To achieve reliable conclusions, it is wise to include a noise penalty determined from rigorous simulations.

Table 1 offers an example of PCIe 5.0 and 6.0 design by loss budget. Included for each is a 4.0 dB penalty budget item for un-simulated noise effects, informed through rigorous simulations results from PCI-SIG work group members. The first two entries represent 1-connector PCB designs for PCIe 5.0 and 6.0. The last column compares the extended reach possible with 7 dB cable assembly budget, achieving 1-meter cable with 3" host and 4" card PCB lengths.

WAYS FORWARD

Satisfying the signal integrity requirements of a receiver is getting significantly harder to achieve. As we continue to push up the Nyquist rate, the loss per inch of the PCB rises significantly. Mid-range PCB dielectrics can have a loss as high as 1.4 dB/in at 16GHz, and breakout region designs around the package and connectors can be around 2dB each. This can very quickly eat away



▲ **Fig. 1** ICN is a weighted, single-number metric that is determined (left) to filter the measured connector crosstalk power sum (right), and thus test the overall connector crosstalk compliance.



▲ **Fig. 2** Samtec's PECFF Emulation Platform showing configurable cabled topologies between mock GPU cards in a typical AI/ML system architecture.

TABLE 1			
	PCIe 5	PCIe 6	PCIe 6 1m Cable
Target	36	32	32
PCB	10" 1.1 dB/Inch @ 16 GHz 11	10" 1.0 dB/Inch @ 16 GHz 11	3" 1.0 dB/Inch @ 16 GHz 3
RC Package	8.5	8.0	8.0
Vias, Caps	1.5	1	1.5
Connector (Spec)	1.5	0.8	
1m Cable & Conns			7
CEM 4" AIC Budget (Ref)	9.5	8.5	8.5
Reflection & Crosstalk Penalty	4	4	4
Total	36	32.3	32.0

the loss budget and significantly minimize component placement flexibility. In addition, crosstalk becomes a greater risk as smaller geometries become antennas at these frequencies, which often leads to more PCB layers to achieve good isolation.

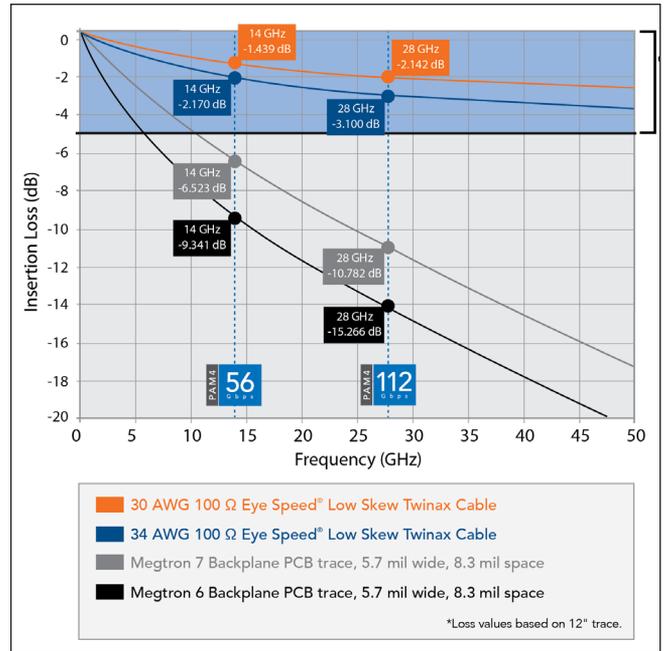
One way to mitigate these challenges is by using a Flyover® cable (see **Figure 3a**) near the transceiver, get-



▲ Fig. 3a Image of Samtec's Double Density Flyover QSFP cable system connecting to Si-Fly® interconnect.

ting the signal off the PCB as soon as possible. Samtec's Flyover cables, for example, are fully shielded differential pairs, so crosstalk is minimized within the assembly. They use high-end uniform dielectrics and a much larger conductor than any high-speed PCB trace, which minimize losses to around 0.15-0.20 dB/in at 16 GHz, depending on the gage used (or 0.18-0.26 dB/inch at 28 GHz, see **Figure 3b**). This allows for much more reach in the system and, therefore, enables placement flexibility.

Without a doubt, the PCIe 6.0 specification is enabling extremely high data rates. Although the standard is backward compatible with PCIe 5.0, designs will have to change dramatically to manage the loss and laten-



▲ Fig. 3b PCB trace vs. Flyover loss, showing 7x cable reach over Megtron 6.

cy issues of the new specification. Although some test and measurement issues may remain, the good news is that interconnect design methodologies and products already exist to get the most out of the latest PCIe 6.0 designs. ■



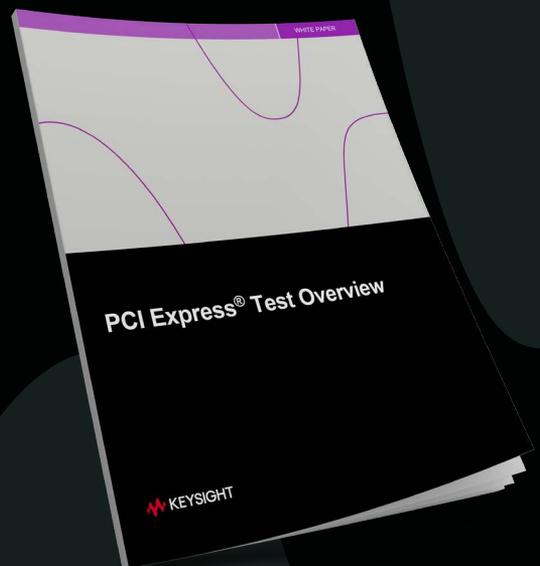
Streamline Your Path to PCIe® 6.0 Success

Achieve PCIe® design success with Keysight's complete and scalable PCIe test solutions — from simulation to protocol.

Innovators
▶ **start here**

Learn more at www.Keysight.com

PCI-SIG®, PCIe® and the PCI Express® are US registered trademarks and/or service marks of PCI-SIG.



Get the white paper

Compute Express Link (CXL)

A New Coherency Protocol Over PCIe

Introduction

Multiple technology trends have combined to fuel the expansion in worldwide datacenter capacity. That capacity has been expanding both in breaking ground on new datacenters, and improving the compute, storage, and network capabilities within and between those datacenters. An important new interface that will be widely deployed in datacenters in the coming years is Compute Express Link® (CXL®), an important technology that leverages the dependability and performance of the PCIe physical layer and introduces new high bandwidth and low latency protocols on top of PCIe for attaching accelerators and memory buffers.

CXL: Meeting the Memory Latency Challenge

For CXL devices to effectively handle the workload demands unique to accelerators and memory buffers, they must operate with much lower latency than what the PCIe interface natively allows. The protocol features and enhancements that are deployed in coming generations of PCIe and CXL present unique challenges to the system designer.

In addition to the latency requirements, it's important to understand how CXL will be used in datacenters to get a full grasp of the protocol challenges. A key characteristic of the modern datacenter is disaggregation and composability. Rather than having compute and storage resources constrained to a single server, new technologies are enabling these resources to be shared across the datacenter, greatly increasing capacity and efficiency. This can reduce the need to overprovision resources to a single server and eliminate islands of resources and utilize them where needed.

A straightforward example of this is the way Non Volatile Memory Express® (NVMe®) reservations and NVMe over Fabrics (NVMe-oF) can be used to share NVMe storage resources over a network and make them appear local, not just semantically, but also in terms of the latency and performance. Sharing storage does have bandwidth and latency needs, but sharing compute resources such as accelerators and memory, and keeping them coherent, is much more difficult since the latency requirements are so much lower.

NVMe drive latencies are typically under 100 microseconds, and when used with certain high performance memory technologies, sometimes referred to as Storage Class Memory or Persistent Memory, can hit latencies near 10 microseconds. While excellent for storage applications, those latencies are still much slower than the sub-20 nanosecond latencies achieved by DDR. CXL is targeting sub-100 nanosecond latencies in order to provide the performance needed for keep disparate memory resources coherent.

Some coherency solutions have managed these challenges through proprietary interfaces, but these have limited adoption and are only applicable to a small set of use cases. The need for an open standards-based protocol for sharing compute and memory resources has driven huge investment in research and development for these memory coherent interconnects, bringing us to CXL. To understand how CXL will enable these coherent memory use cases, its necessary to understand how CXL uses the technology underneath, PCIe, and the validation requirements.

PCIe Progression

When we look at how PCIe is deployed in the datacenter, it's clear that there is a dual transition happening. On the one hand, we see the migration from PCIe 4.0 at 16 GT/s to PCIe 5.0 at 32 GT/s to PCIe 6.0 at 64 GT/s. In particular, the step from PCIe 5.0 to 6.0 will introduce challenges to test and validation in both the electrical and protocol layers. There is a doubling of bandwidth in PCIe 6.0 with the introduction of PAM-4 signaling. PAM-4 signaling encodes 2 bits in each electrical transition on the wire. This is achieved by 4 level signaling, which can be more susceptible to bit errors. The PCIe protocol will compensate for this by introducing Flit encoding and Forward Error Correction. With multiple changes happening with the transition from PCIe 5.0 to 6.0, a comprehensive protocol validation plan will be necessary. Keysight is focusing resources on providing a comprehensive suite of solutions for PCIe 5.0 and 6.0 for TX, RX, and Protocol needs.

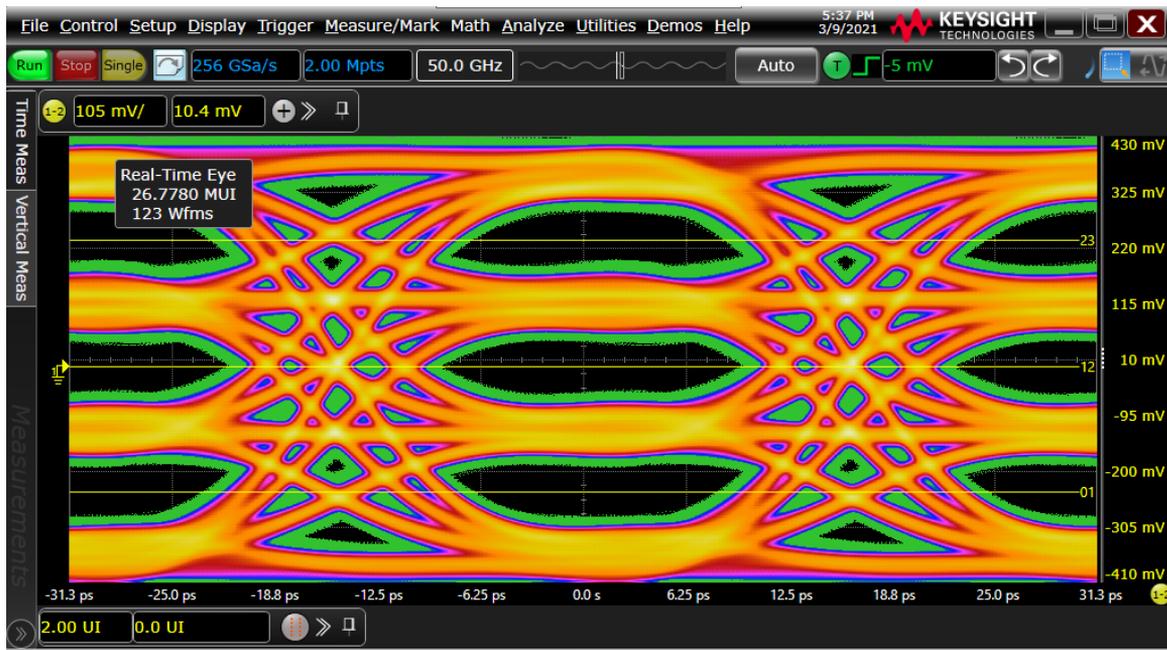


Figure 1. PAM-4 Signaling introduces new complexities.

In parallel with these fundamental changes to PCIe operation, we see the PCIe interconnect being leveraged for new use cases by the CXL protocol, enabling the connection of and sharing of compute and memory resources over a high bandwidth and low latency interconnect. While the electrical subblock for CXL devices will be the same as for PCIe device, the link layer operation is different. During Link Training, an Endpoint device can indicate support for CXL via the Alternate Protocol Negotiation. Based on that the host will access the endpoint either using regular PCIe operations or CXL operations. This is referred to in the CXL specification as Flex Bus operation and introduces the first of several protocol validation challenges.

The CXL Flex Bus

The CXL specification defines how a slot in a host system can operate as either a PCIe slot or a CXL slot. This is referred to as the Flex Bus. During link training, CXL devices start the negotiation normally at 2.5 GT/s using 8b/10b link training. Then, during the configuration state of the Link Training State Machine (LTSSM) the CXL device can indicate that it supports CXL during the alternate mode negotiation. Assuming both sides of the link support the alternate mode, link training progresses according to the PCIe specification for LTSSM by entering the L0 state at 2.5 GT/s. It's important to note that CXL operation requires a link speed of at least 8 GT/s, although 8 GT/s is considered a degraded mode of operation. If CXL mode is selected during link training, the CXL specification requires that both sides will support at least 32 GT/s operation. The link speed can be negotiated down to 8 GT/s, but if the link is unable to negotiate a link speed of 8 GT/s or higher, the link will not operate. Normal operation for CXL uses link speed of 32 GT/s or higher for the necessary performance.

These differences in CXL operation relative to PCIe operation can introduce difficulties for validation engineers. While the PCIe LTSSM state machine may be well understood, many engineers are not as familiar with the Alternate Mode Negotiation necessary to enable CXL operation. Tools that provide a clear view of the different phases with the Configuration state of LTSSM will be necessary to understand whether Alternate Mode Negotiation is executing correctly or not. By leveraging the superior signal integrity enabled by the unique design of their protocol analysis solution, Keysight's protocol tools give an accurate understanding of the operation of the LTSSM state machine.

Since the Flex Bus can negotiate to either CXL operation or regular PCIe operation the system designer is granted much more freedom since it's not necessary to predict the exact number of CXL lanes versus PCIe lanes needed at the time of system design. Of course, different use cases will require a different number of lanes, and designers will need to be alert to expected applications to allocate the correct number of lanes to each slot in the system, but can be confident that the slot will be able to support any PCIe or CXL device depending on the use case of the system. While the electrical signaling is the same, the Link and Transaction layer are markedly different in CXL when compared to PCIe. These optimizations for CXL introduce a certain set of assumptions about the transaction that allow for lower latency and overhead when compared to normal PCIe operation. These optimizations are critical for CXLs intended use: low latency transaction to cache and memory. One critical optimization is the use of Flow Control Unit, or Flit, encoding.

Flow Control Unit (Flit)

In normal PCIe operation with PCIe 5.0, there's a variable payload. Given the wide variety of PCIe use cases, enabling a variable payload size makes sense for versatility. However, the framing that needs to be added to each transaction to manage and track this variable payload adds latency. Therefore, CXL works with a fixed payload size of the Flow Control Unit (Flit). The fixed flit size enables several assumptions and optimizations to be built into the CXL protocol which reduce overhead and latency. It's worth noting that PCIe 6.0 introduces its own Flit mode to alleviate the detrimental latency effects of Forward Error Correction (FEC), however the PCIe 6.0 Flit format is different than the CXL Flit format.

The CXL flit uses a fixed 528b length which includes a 16b CRC for error detection. All flits are preceded by a 16b protocol ID to indicate which sub protocol the flit is for: CXL.io, CXL.memory, CXL.cache. The fixed size for flits enables simpler retry mechanism which doesn't require sequence numbers to be passed back and forth between the host and device, eating up overhead. This necessitates a retry buffer which contains more flits than the round-trip delay between the host and device, but significantly reduces latency in the event of a retry. With such a fundamental change to the lower layers of the interface, it's easy to see how critical it will be to fully validate the protocol operation and performance of the CXL Link and Transaction layers with trusted analysis capability. Especially with respect to retry operations which are context sensitive, and proper debugging of retries will require information that is not present in the flits being retried. Tools capable of tracking and maintaining expected sequence numbers will be necessary for proper debugging.

Sub-protocols and their Use Cases

Not all CXL product types will need all CXL protocol features. In order to minimize the complexity of designs with specific use cases, CXL is deployed as 3 sub-protocols: CXL.io, CXL.memory, and CXL.cache. An examination of each subprotocol and its capabilities will help to gain understanding of what sorts of devices are possible with each sub protocol.

CXL.io, which is nearly identical to the PCIe transaction layer, and supports basic functionality like discovery, configuration, and interrupts. All CXL devices are required to implement CXL.io.

CXL.cache allows attached accelerators to access the CPU attached memory and ensure that the device cache is coherent. This optional protocol will be implemented by accelerators, referred to as Type 1 devices in the CXL specification. The accelerator may be a special type of processor optimized for a particular kind of computation. The CXL.cache protocol will enable the accelerator to access the CPU attached memory, ensure that the devices onboard cache is coherent, and use that coherent data with minimal overhead. This can enable computations to be done more quickly and can also reduce the usage of a potentially power hungry or busy CPU.

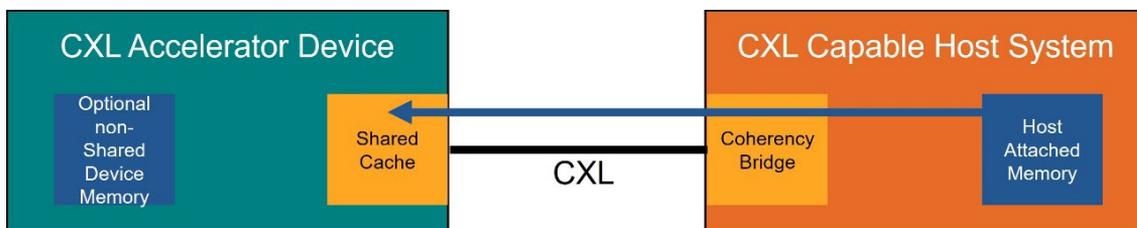


Figure 2. CXL.cache allows sharing cache between a host and acceleration device

CXL.memory is an optional protocol which enables a CPU to access memory in a memory expansion device or buffer, referred to as Type 3 devices in the CXL specification. A key use case here is the so-called persistent memory, which slots between DRAM and NAND Flash in terms of access latency. These Type 3 devices will enable expansion of memory in a system via the PCIe bus. This enables all sorts of trade offs and options depending on cost and performance needs.

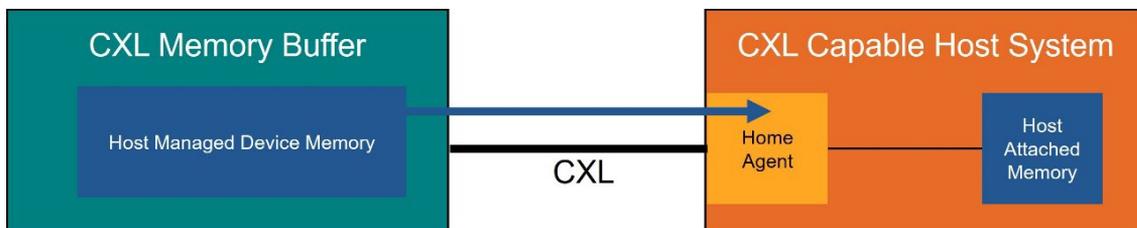


Figure 3. CXL.mem allows a host to access memory on an attached Memory Buffer device

Accelerators with Memory on board, referred to as Type 2 devices in the CXL specification, will implement both CXL.cache and CXL.memory. Implementing both the CXL.cache and CXL.memory protocols independently allows an endpoint to access to the memory of a CPU and the CPU to access the memory of the attached endpoint, allowing the sharing of memory resources in both directions.

By dividing the capability between these sub protocols, complexity is reduced where certain features are not needed. For example, simple memory expansion buffer does not have the responsibility to manage coherency, it just needs to be accessible by the CPU, so it does not need to implement the CXL.cache protocol.

While utilizing optional sub protocols can reduce the complexity in some designs, it increases the complexity of analysis and validation. Validation engineers will need tools that can enable them to identify the sub protocols, but also check the behavior of those protocols on the link.

Coherency and Biasing

In a multi-processor system, such as a CPU with one or more accelerators attached, having coherent cache memory contents has many benefits. For example, the CPU can manage which attached accelerators do what computation on particular operands within the cache in order to execute a particular operation in the most efficient manner, whether the goal be power savings, speed, or both.

However, maintaining coherency among those disparate caches has costs in terms of overhead. In many cache coherent protocols, there is the concept of the Snoop operation. Typically, this operation is intended to notify other agents in the cache coherent network, that the remote shared cache that they are working with has been updated, and that they need to update their local cache. It's easy to see that very quickly, too many snoop operations could cause a flurry of data copying, incurring a great deal of overhead, and ultimately overriding the benefits of having multiple processors working on an operation.

There are many theories of operation that attempt to balance the need to maintain coherency while reducing the need for data copying and its associated overhead. One strategy leveraged in the CXL specification is the Bias Based coherency model. In this model a device-attached memory can be Device Biased or Host Biased.

When the device-attached memory is in the Device Bias state, the device can access that memory without checking with the host via request or snoop operations. This is possible because by virtue of that device-attached memory being in the Device Bias state, the device knows that the host does not have that line in cache, and the data in that cache is not stale. Therefore, there's no need to check with the host before operating on that data. This saves unnecessary back and forth overhead.

When the device-attached memory is in Host Bias state, the device must resolve coherency for a particular line of cache before accessing it. This requires snoops to the host. This ensures coherency but has associated overhead.

By enabling the system to toggle the bias of device-attached memory between Host Bias or Device Bias, the system can ensure coherency with minimal overhead. In the Host Bias state, the overhead of snooping and updating cache lines is incurred where it is needed. In the Device Bias state, that overhead is eliminated, because it is not needed.

Ensuring that a system is properly following the rules for biasing, has profound implications for the accuracy and coherency within a system, as well as its performance. Improper behavior around biasing could be easily masked, as the system may function properly. Memory access would be possible, and coherency maintained, but there could be unnecessary overhead if biasing rules are not followed properly. Analyzing and detecting improper behavior around biasing can yield important insights to improve system performance and reduce latency

CXL Progression

Initial devices will follow the CXL 1.1 specification with CXL 2.0 devices to follow. While CXL 2.0 devices will be backwards compatible to CXL 1.1 devices it's important to understand key features that are introduced in CXL 2.0. CXL 1.1 was designed with the aim of being a coherent memory interconnect within a single node. In CXL 2.0 the addition of switching enables multiple nodes inside a rack or chassis to share and access memory. Via CXL enabled switches, the number of devices that can be attached can be greatly increased. Alongside switching, CXL 2.0 enables pooling of resources. CXL Pooling is the ability to allocate and deallocate resources in different physical locations to different hosts. This brings the system implementation closer to a true memory disaggregated system where memory resources can be fully shared among different applications.

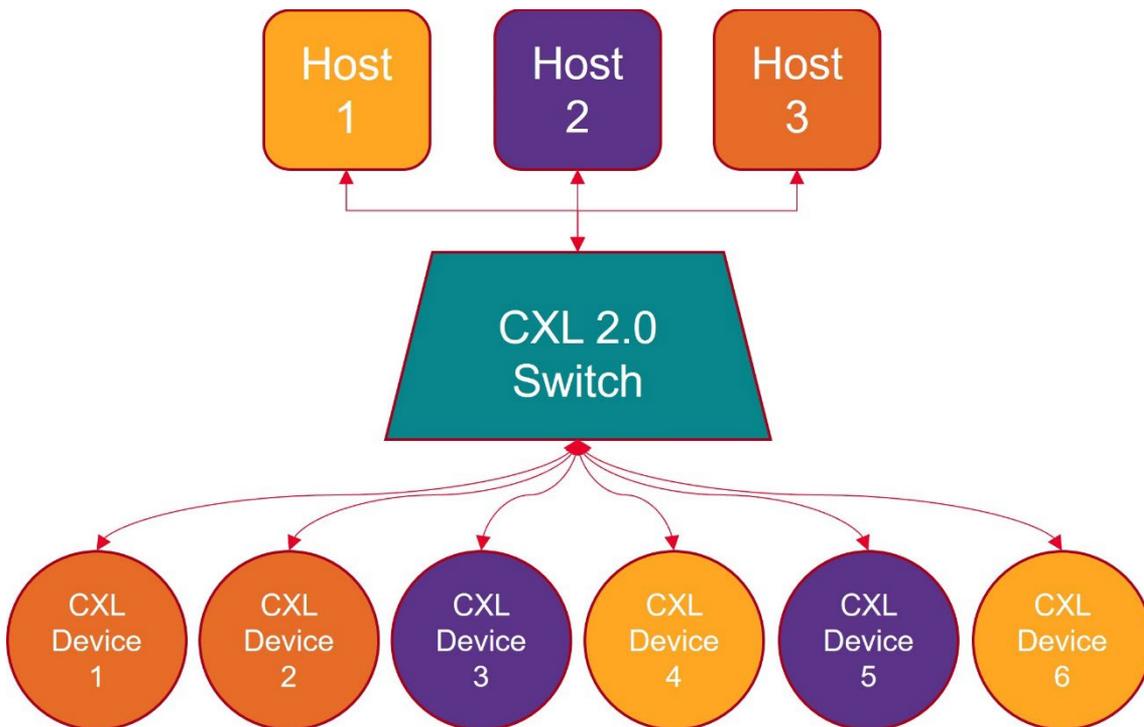


Figure 4. CXL switching introduced in CXL 2.0 allows the pooling and sharing of CXL resources

What is Keysight doing?

Keysight is working closely with key partners to deploy solutions for CXL, PCIe 5.0 and PCIe 6.0. This portfolio of solutions enables deep insights and Keysight is committed to supporting the innovation required to keep that capacity expanding with a broad portfolio of test and measurement tools that cover all aspects of the datacenter. Several important datacenter interfaces are in transition such as 800G Ethernet, DDR5, and PCIe 6.0 and Keysight is actively supporting these transitions with our leading physical layer measurement and protocol analysis and generation tools.

Keysight enables innovators to push the boundaries of engineering by quickly solving design, emulation, and test challenges to create the best product experiences. Start your innovation journey at www.keysight.com.



This information is subject to change without notice. © Keysight Technologies, 2018 – 2022, Published in USA, December 13, 2022, 3122-1353.EN



Prepare for PCIe® 6.0

Understand how the PCIe® standard has evolved and how to comply — regardless of generation.

[Read white paper](#)

**Innovators
▶ start here**

Learn more at
www.Keysight.com

PCI-SIG®, PCIe® and the PCI Express® are US registered trademarks and/or service marks of PCI-SIG.

USB4 Version 2.0 from Simulation to Tx, Rx and Interconnect Test

Keysight

WHAT IS USB4 VERSION 2.0

The USB-IF released the USB4 Version 2.0 Specification in October 2022.

With this latest specification, each link will have four bi-directional differential lanes, and each lane will run at 25.6GBaud, 40Gbps. In the symmetric mode, each link will have 2 lanes running at 40Gbps for an aggregate 80Gbps in each direction. With a new asymmetric mode, the link can be negotiated to transmit 3 lanes in one direction. The net result is 120Gbps in one direction, and 40Gbps in the other direction in asymmetric mode.

To increase and double bandwidth, this next-generation USB technology chooses PAM3 at 25.6GBaud and 40Gbps using 11bit/17trit encoding. This small increase in the fundamental frequency from 10GHz to 12.8GHz allowed the use of existing USB4 and Thunderbolt 4 cables and connectors.

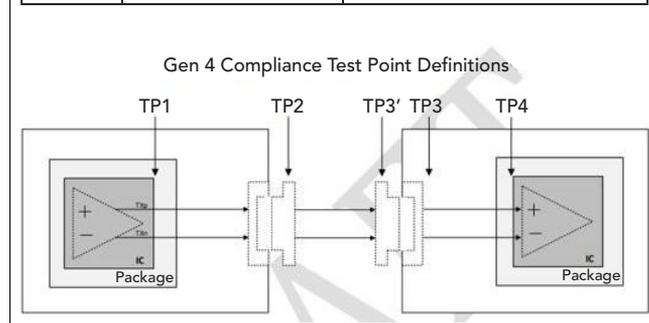
For PHY layer electrical validation engineers, each of the 4 differential pairs will run at 25.6GBaud, PAM3, and 40Gbps.

USB4 VERSION 2.0 SIMULATION

The cost of turning silicon is very high plus it adds delays to a first to market strategy. A key component of reducing this risk, and increasing success is to perform rigorous end- end simulation of the entire USB 80Gbps link. To enable early design stage simulations as well as extensive system level post-layout analysis, simulation solutions for USB4 Version 2.0 that incorporate IBIS-AMI model makers facilitate the development of such models for USB 80Gbps devices. These IBIS-AMI models are

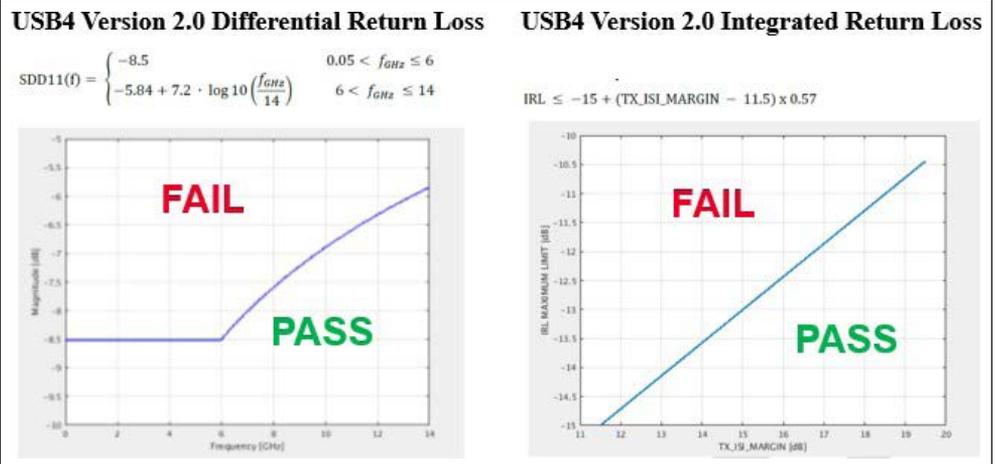
then used in channel simulations to predict and simulate the BER, eye metrics, and other design parameters.

Test Point	Description	Comments
TP1	Transmitter IC Output	Not used for Gen 4 electrical testing.
TP2	Transmitter Port Connector Output	Defined at the output of a compliance plug fixture.
TP3	Receiver Port Connector Output	Defined at the receptacle side of the connector. All measurements at this point shall be done while applying the reference equalization function.
TP3'	Receiver Port Connector Input	Defined at the output of a compliance plug fixture.
TP4	Receiver IC Input	Not used for Gen 4 electrical testing.



▲ Fig. 1 USB4 Gen4 compliance test point definitions.

With USB 80Gbps PAM3, signal and power integrity are critical as frequency and speed increase in printed circuit boards (PCBs). Losses associated with transmission line effects will cause failures in Gen 4 devices. It will be crucial to model traces, vias, and interconnects to simulate the board accurately. Plus improve high-speed link performance in PCB designs with integrated circuit design and electromagnetic simulators customized for power and signal integrity analysis.



▲ Fig. 2 USB4 Version 2.0 TX and RX Return Loss test limit requirement.

USB4 VERSION 2.0 TRANSMITTER TESTING FUNDAMENTALS

To ensure precise TX characterization, there are a couple of key Test Points (TP) that need to be well understood.

TP2 is defined at the output of the compliance test fixture and where all TX measurements are performed. What's different from Gen 2/3 is there are no longer measurements performed at TP3 for Gen 4.

However, a captive/tethered cable device measurement is always performed at TP3 because the only accessible measurement point is at TP3.

There are the traditional jitter and timing measurements from Gen 2/3 but now performed on a PAM3 signal.

As mentioned above, the most significant challenge with Gen 4 is the vertical signal margin. Hence there are a number of new vertical measurements introduced with Gen 4 to ensure acceptable BER performance.

Level_Mismatch compares the difference in eye opening between the "top" and "bottom" PAM3 eye.

Using LFPR (Linear Fit Pulse Response), sigma n, and sigma e, SNDR (Signal Noise Distortion Ratio) is another key vertical margin parameter.

Integrated Return Loss (IRL) takes into account the signal quality at the transmitter coupled with the Insertion Loss (IL) of the channel for another important vertical margin parameter.

Gen 2/3 had 16 TX Presets. To no surprise, Gen 4 with more loss has 42 Presets. Just like Gen 2/3, a sweep needs to be performed on all 42 Presets to characterize each individually, and also to determine the optimal Preset for minimum DDJ (Data Dependent Jitter).

If a USB 80Gbps product is anything larger than a USB stick, there is a high likelihood a re-timer is required. This requires clock-switch measurement like the frequency variation training measurement.

USB4 VERSION 2.0 RECEIVER TESTING FUNDAMENTALS

Like Gen 2/3, Gen 4 has the typical stress cocktail components like launch voltage, ACCM, PJ, RJ, and loss channel.

Like Gen 4 TX, Gen 4 Rx introduces new vertical stress components like SNDR, Level Mismatch, DMSI, and CMSI.

Like Gen 2/3, there are the 2 use/test cases:

- Short channel for low loss products or link partners connected with active cable.
- Long channel for link partners connected with a passive cable.

Gen 4 added a 3rd use/test case for link partners connected to a Linear Re-driver (LRD) cable.

Similar to Gen 4 TX, there is also a special test case for tethered/captive devices. The draft version of the CTS requires that aggressors be added for RX testing.

ADDITIONAL TEST REQUIREMENTS FOR USB4 VERSION 2.0

Similar to Gen2/3, Side-Band (SB) testing will be required to ensure proper link negotiation. New requirements for TX and RX LFPS testing are planned to be introduced with Gen 4. As mentioned above, Gen 4 added an asymmetric operation mode, that also requires specialized testing.

It cannot be stressed enough that vertical margin will be challenging. Hence, the test methods require noise compensation to reduce the effects of measurement noise in the test system.

USB4 VERSION 2.0 TRANSMITTER AND RECEIVER RETURN LOSS

USB4 Gen 2/3 introduced new differential and common mode return loss requirement for Transmitter and Receiver. This is carried on in Gen 4. Return loss, the ratio of reflective power to incident power is the direct measure of the impedance match of a transmis-

sion line. Meeting the test limit requirement is critical to meet compliance certification and ensuring product performance and interoperability.

In USB4 Version 2.0, the Integrated Return Loss (IRL) measurement was introduced, and the differential return loss requirement remains one of the informative items for compliance certification. The increased bit rates to 80Gbps to support USB4 Version 2.0 protocol introduce additional signal integrity challenges and require a more stringent integrated (summed) return loss test corresponding to the integrated power spectral density of the incidental/reflective behaviors over the baseline baud rate frequency range.

The transmitter Integrated Return-Loss is extracted as follows:

$$IRL = 20 \cdot \log_{10} \left(\sqrt{\frac{\int_0^{20GHz} |V_{in}(f)|^2 \cdot |S_{dd22}(f)|^2 df}{\int_0^{20GHz} |V_{in}(f)|^2 df}} \right)$$

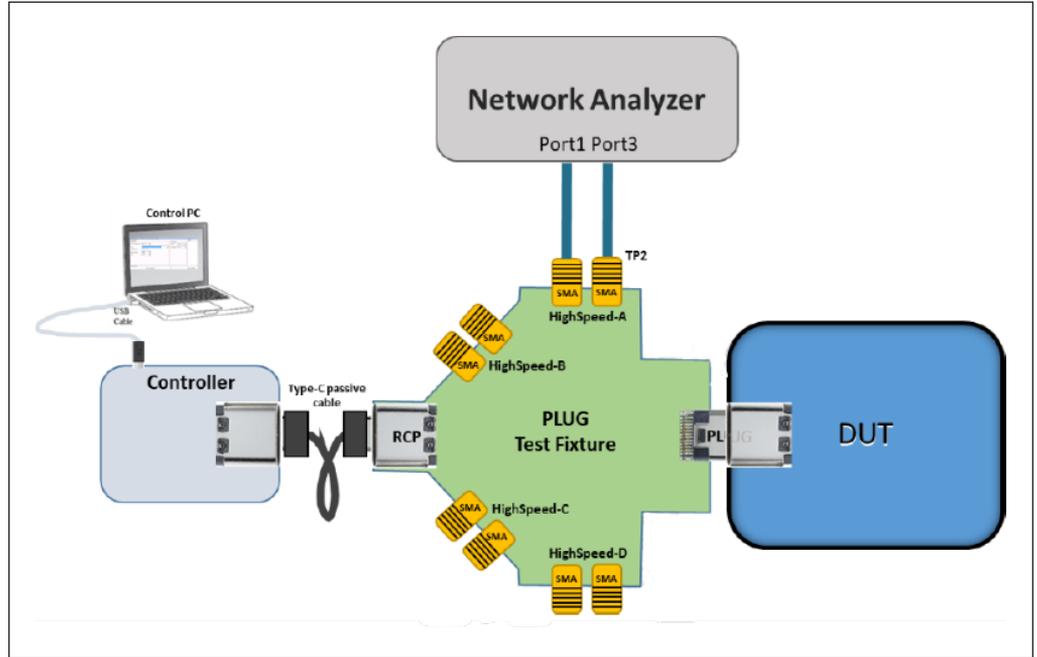
where:

- Sdd22(f) is the Return-Loss of the transmitter at TP2, referenced to single-ended load impedance of 42.5Ω.
- $V_{in}(f)$ is the spectrum of the ideal PAM signal with a 20% slew rate, defined as
- $V_{in}(f) = \sin(\pi \cdot f \cdot T_r) / \pi \cdot f \cdot T_r * \sin(\pi \cdot f \cdot T_b) / (\pi \cdot f \cdot T_b)$, with $T_b = 39.0625ps$ and $T_r = 0.2 \cdot T_b$

The Transmitter and receiver differential return loss (Sdd22) and integrated return loss (IRL) will be setup as below. An S-Parameter touchstone file (S2P) will be measured by a Vector Network Analyzer (VNA) while the transmitter/receiver DUT is in active mode with PRBS7 pattern driving by a USB4 test microcontroller and ETT tool. The measured S-Parameter will be analyzed by the SigTest tool for delivering test results.

The transmitter IRL maximum limit is a function of the measured transmitter ISI margin (TX_ISI_MARGIN) which corresponds to the transmitter signal-to-residual ISI ratio.

Therefore, to verify the verdict of the transmitter IRL, a waveform file (.bin) of ui_jitter_vertical that is used in the transmitter's Timing and Voltage measurement test will be compiled using the USB4 SigTest tool.



▲ Fig. 3 USB4 Version 2.0 TX and RX Return Loss proposed test setup.

As discussed, in the transmitter and receiver return loss test it may be possible in some cases for the TX or RX signaling of the DUT to introduce error into the measurement. Improperly configured DUTs (i.e., the DUT is not transmitting the proper pattern, or being forced into the incorrect mode, etc.) can produce erroneous measurement results, which could lead to false failures. Care should be taken to ensure that all sources of error have been minimized. Besides, a VNA used for the measurement should necessarily warm up and be calibrated with an Electronic Calibration Module (ECal) prior to the measurement. The VNA's setup should follow proposed CTS requirements with Port 1 and Port 3 connected to the DUT with a sweep frequency range from 50MHz to 20GHz and IFBW at least 1600 points to minimize the trace noise.

NEW TEST REQUIREMENTS FOR USB4 GEN3 AND GEN4 CABLES

Compared to the previous USB 3.2 Type-C CTS, USB Type-C CTS for USB4 is much more complex. The increased bit rates to 40Gbps/80Gbps to support USB4 and USB4 Version 2.0 protocol introduce additional signal integrity challenges and require more stringent integrated test parameters corresponding to the incidental/reflective behaviors over a frequency range.

The new test group – Test Group B-8 and Test Group A-8 requirements are aimed at the integrated S-parameters (except for insertion loss and differential-to-common-mode conversion) to avoid the potential rejection of a functioning cable assembly that may fail the traditional S-parameters spec at a few frequencies. In the case of integrated return loss (IRL), it now manages the

reflection between the cable assembly and the rest of the system (host and device) with more IRL allowed if the cable loss is smaller.

SUMMARY

USB Gen 2/3 at 10/20Gbps NRZ continues to be very challenging to implement.

USB Gen 4 at 25.6 GBaud PAM3, 40Gbps 11bit/7trit encoding, coupled with asymmetric mode crosstalk, and the same loss channel makes it exponentially more complicated.

The IP/PHY development kits, simulation solutions, and T&M solutions reviewed in this paper are the foundational blocks required for silicon and system integrators to design compliant USB4 Version 2.0 products that meet strict interoperability requirements in the vast Type-C ecosystem. ■



Boost the Potential of Your Type-C Designs

Get step-by-step procedures for debugging and optimizing the USB link.

[Read application note](#)

Innovators
▶ **start here**

Learn more at
www.Keysight.com

USB-C Signal Analysis Methodology

Brian Holden
Kandou Bus

USB 3.x systems frequently use redriver devices located near the USB connector to clean up the signals. This allows impairments caused by the traces within the system to be hidden from the externally facing interface. The higher speed of USB4 standard has prompted a move to retimers and away from redrivers within USB-C systems for proper signal integrity.

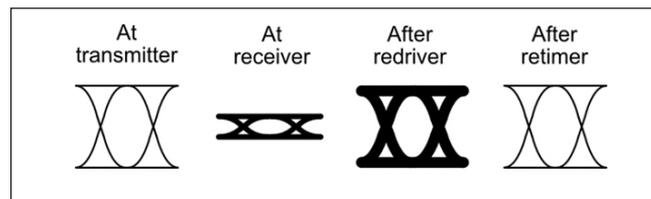
Redriver devices are limiting analog amplifiers and fully recover the bit stream by using equalization and clock-data recovery. They then retransmit a fresh copy of the recovered data stream. Retimers reduce the jitter and retransmit the signal. **Figure 1** shows a comparison of representative transmit-direction eye diagram that have been sent through a redriver versus a retimer.

Jitter accumulation is the key factor that makes redrivers fail in real systems at higher speeds. To analyze a system design to see whether it will function well and meets the requirements of USB4, detailed signal integrity analysis must be undertaken.

The following steps can be taken:

Collect a set of S-parameters for each of various representative cables that will be connected to the USB-C connector. These can either come from measurement by a network analyzer or from the cable manufacturer or both. These sets should contain a mixture of losses. Representative losses for passive cables include 2.5, 5, 8 and 10 dB. These roughly map to 0.2, 0.4, 0.8 and 1 meter in length. Your collection should also include cables that are at the edge of the specified parameters for skew and crosstalk. There should be a cable that is at the extreme bounds of all parameters at the same time. It is also important to include some cables that violate the specifications. Twelve-port S-parameter files are useful so that the crosstalk from two aggressors can

be accounted for in the measurements. Retimers contained in active cables are treated as full endpoints. Active cables are modeled as concatenated full links with passive segments in between.



▲ Fig. 1 A comparison of redriver-retimer signal.

Collect a set of S-parameters for each of a wide variety of far-end systems. These systems should have a mix of trace lengths between the USB-C connectors and devices. These systems should have imperfections in terms of skew and crosstalk, both at and beyond the bounds of the specifications. Chip vendors often provide two-port IBIS AMI models of their devices and sometimes of reference platforms. If it is possible to get a bare board, the board's S-parameters could also be measured by a network analyzer. These systems may have a mixture of redrivers and retimers. If retimers are used in a system, model it as a concatenation of full links. If retimers are used, the portion of the system behind the retimer can often be ignored because the link between the retimer and the endpoint chip may have plenty of margin.

Collect a set of S-parameters for the near-end system to be analyzed. This should accurately model each USB-C port separately. Some combination of vendor measurements and your measurements may be necessary. Any redrivers or retimers should be accurately modeled

as detailed below. The segments before and after these devices should be modeled including effects such as crosstalk and skew. Again, it may be possible to ignore the link between a retimer and processor if this link is found to have plenty of margin.

The S-parameter sets for the various elements should be collected with care, using the same maximum frequency. Measurement effects should be carefully de-embedded. Any extrapolations should also be applied with great care.

Prepare a simulation model of the source and destination devices. On the transmitter side, this model should accurately account for the transmit swing, the transmit Signal to Noise and Distortion Ratio (SNDR), and the transmit package. It should have a method of setting the transmit equalizers automatically according to the channel used with some defined inaccuracy. On the receive side, this model should accurately account for the receive package, the input noise, the equalizers, and the sampling uncertainty caused by input jitter & by the operation of the clock-data recovery circuit. The model should have a method of setting the receive equalizers automatically according to the channel used with some defined inaccuracy.

A redriver model should be prepared. That model should have a gain element, an input noise source, and a transmit distortion model. This model should be correlated with the performance of actual devices. More sophisticated models can incorporate additional effects such as power supply noise impacts.

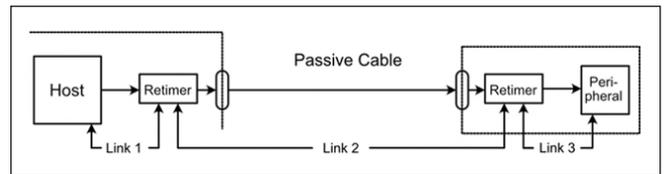
Next a retimer model should be prepared. The most straight-forward models look like the source and destination devices mounted in smaller packages. More sophisticated models can incorporate subtle effects such as the impacts on the reference frequency when passed through chains of phase locked loops.

Purchase one of the commercially available analysis tools and see that personnel are trained and practiced in their use. We use Keysight ADS but other tools are fine as well. There is a long learning curve for these tools. The first series of simulations performed by any new group of engineers will inevitably be inaccurate due to missteps both small and large.

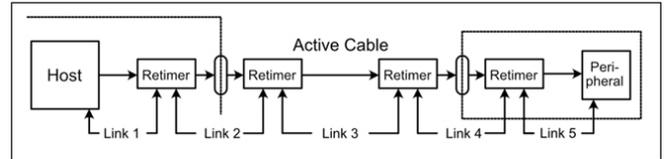
For each element individually, a series of specific compliance measurements must be taken at specific points as defined by the USB specifications and by other definitions of interest. These include items such as the total insertion loss for the host and the integrated return loss for cables. Given that USB4 and its compliance methodology is still emerging, it is reasonable to expect some flux on the details of these measurements.

Once the elements are measured, the performance of the system must be looked after in a more holistic manner than is defined by the specifications. Determine the target Bit Error Ratio (BER) that the system should operate at. Determine a suitable correction factor to add to the analysis to cover the collective negative impacts of small effects that are not otherwise covered.

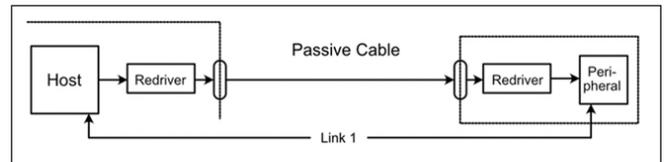
Run the tools and determine the margin for each collection of source device, source traces, source redriver/retimer, further source trace & connector, passive/active cable, destination connector and trace, destination



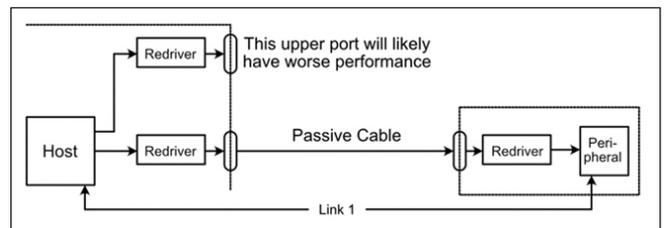
▲ Fig. 2 Retimers break the full link into three shorter links.



▲ Fig. 3 With active cables, the full link is broken up into five shorter links.



▲ Fig. 4 Redrivers leave the link subject to the concatenation of impairments.



▲ Fig. 5 With redrivers, the USB-C port further from the host often has poorer performance.

redriver/retimer, further trace, and destination device. Both directions must be simulated. This analysis will typically involve hundreds of runs and may take many hours of run-time. Find the combinations where the remaining margin becomes negative. Hopefully, these will not include some of the desired combinations. When retimers are used, this simulation burden is much reduced because the full link is broken up into a sequence of shorter links as shown in **Figure 2** and **3**.

Redrivers are known to be problematic for USB4 because they do nothing to remove the jitter in the signal and drive out a noisier signal than retimers do as shown in **Figure 4**. The transmit jitter specification is typically the first to fail in these cases

For redrivers, the impact of all this analysis complexity presents another compelling reason to prefer the use of retimers. The links on the system sides of the retimers can often be ignored if they have plenty of margin, significantly simplifying the design.

Adding to the complexity, different ports on the same system are likely to have different performance for end-users when redrivers are used. For example, a port located in a far corner of a system may experience more errors than a port located in the middle as shown in **Figure 5**. If retimers are used near each USB-C port, the system will be much simpler to analyze and will have more consistent operation over a wider variety of use cases.

Once systems are built, compliance at the USB-C connector can be measured in a variety of ways. Measurement is important to help control the manufacturing variance of systems and to debug problems in the lab or in the field. It is also important to help correlate the analysis techniques used to measured values.

A first method of measuring compliance is to build a compliance board that interfaces the USB-C to high-speed connectors over short, constant length traces. Coax cables are then used to connect these to a high-speed oscilloscope. Care must be taken to de-embed the board and the cables from the measurement.

A second method of measuring compliance is to create a paddle card that has a retimer device on it that contains eye-scope functionality. This retimer device can be placed very close to the USB-C connector to help minimize the need for de-embedding. Compliance tests can then be constructed using the eye scope and software. Vertical and horizontal bathtub curves can also be constructed. Internal eye-scopes can also be useful in the manufacturing test of both the devices and the systems that use them. The Matterhorn retimer from Kandou contains such an eye-scope.■

About Brian Holden

Brian Holden is vice president of standards at Kandou. Previously, he served as president and fellow of the HyperTransport Consortium, chair of the market awareness and education committee for the Optical Internetworking Forum (OIF) and director of standards for PMC-Sierra (now Microsemi). Brian began his career as an electrical engineer at GTE. He received a Bachelor of Science degree in Electrical Engineering from the University of California, Davis and an MBA from Cornell. He is the author of "HyperTransport 3.1 Interconnect," a biography of his great grandfather titled, "Charles W. Woodworth: The Remarkable Life of U.C.'s First Entomologist," and the forthcoming book titled "Chord Signaling." He has 49 U.S. Patents.

DDR5 Signal Integrity Fundamentals

Tim Wang Lee
Keysight

In July 2020, a new standard for double data rate (DDR) memory was announced. The exciting DDR5 technology promises higher data rate with reduced power consumption. This is a promise that is familiar to serial link designers. However, like most things in life, there is no free lunch. The advances for lower power and higher speed come with an increase in design complexity. The most notable difference between DDR5 and previous generations is the introduction of decision feedback equalization, which is a technique used in serial link systems to improve the integrity of received signals.

In the wake of the new technology, this article examines some fundamental signal integrity concepts in the context of DDR5. The first section introduces the eye diagram: a metric to determine the goodness of signal integrity. The second section describes root causes of signal integrity problems by examining the single pulse response. The third section prescribes possible solutions to the resulting signal integrity problems.

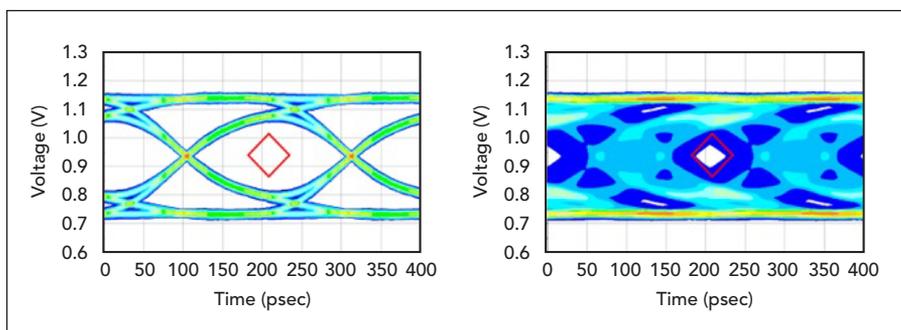
EYE DIAGRAM TO DETERMINE SIGNAL INTEGRITY

The eye diagram is a primary metric for evaluating the signal integrity of a channel. It is created by appropriate processing of received pseudo random binary sequence (PRBS) through a channel. To create an eye diagram in the context of the “write” cycle of memory operation, the controller (transmitter) sends PRBS through a channel to reach the

memory module (receiver). The received PRBS pattern at the memory module is divided into segments with the same time interval. These segments with identical time interval are then stacked on top of each other to create an eye diagram.

In **Figure 1**, there are two eye diagrams in blue and eye masks in red. By comparing the eye diagram at the output of the channel to an eye mask, one determines the signal integrity of the channel. The eye mask is a graphical representation of the receiver’s threshold. The eye mask shows the acceptable timing and amplitude of the received signal for a given bit error ratio (BER).

As shown on the left of Figure 1, the eye is open. The channel has good signal integrity when there is no overlap between the output eye diagram and the eye mask. If the output eye diagram does not overlap with the eye mask, the receiver can determine a digital one or digital zero based on the received analog voltage level and timing. On the other hand, if there is an eye mask violation



▲ Fig. 1 Left: the eye is open because there is no eye mask violation. Right: the eye is closed because of the eye mask violation.

tion (as shown on the right of Figure 1) the eye is closed. A digital one or digital zero cannot be distinguished at the receiver. The eye diagram gives engineers a metric for the performance of a given channel. When there is a closed eye at the receiver, one needs additional analysis techniques to identify the root cause of the eye closure.

FREQUENCY-DEPENDENT LOSS AND REFLECTION IN DDR5

The main concerns specified in DDR5 standard are reflections and frequency-dependent loss.¹ **Figure 2** shows the single pulse response of a DQ line from the controller to the memory module. The single pulse response is the received waveform at the memory module when a single pulse, a digital one, is sent from the controller.

In Figure 2, the red dotted line is the ideal case where no reflections or frequency-dependent loss are in the channel. In blue, one observes the frequency-dependent loss of the channel as the spreading of the ideal pulse. The reflections in the channel come later in time. Because the spreading of the single pulse and the reflections can interfere with other pulses, one often refers to them as inter-symbol interference (ISI).²

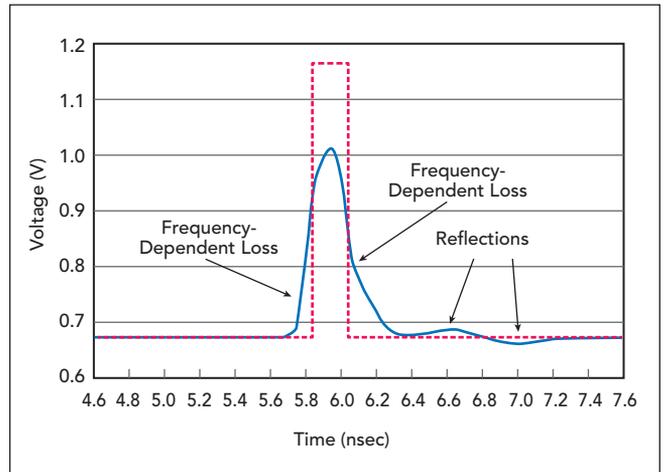
The ISI caused by frequency-dependent loss is common in serial link channels while the reflection problem caused by impedance discontinuities is quite unique to DDR.

DECISION FEEDBACK EQUALIZATION IN DDR5

If the root cause of the signal integrity problem is frequency-dependent loss, the most straight-forward solution is to reduce the length of the channel or use low loss material in fabrication. To minimize the amount of reflections, traces should be designed with controlled impedance. If the eye remained closed with appropriate channel length, fabrication material, and impedance control, equalization at the receiver can help further improve/open the eye at the receiver.

In DDR5, four-tap decision feedback equalization (DFE) was specified to mitigate the loss and reflection without amplifying noise.¹ With each tap representing one unit-interval, the four-tap DFE corrects up to four unit-intervals after the current received bit. As the name suggests, the decision feedback equalization algorithm makes a decision on each received bit and feeds a modified version of the bit back to the receiver.

In the DFE algorithm, the received analog waveform first arrives at the symbol detector. The symbol detector decides whether the received analog waveform represents a digital one or a digital zero. If the detected



▲ **Fig. 2** Single pulse response of a DQ line shows both frequency-dependent loss and reflections in a DDR5 channel.

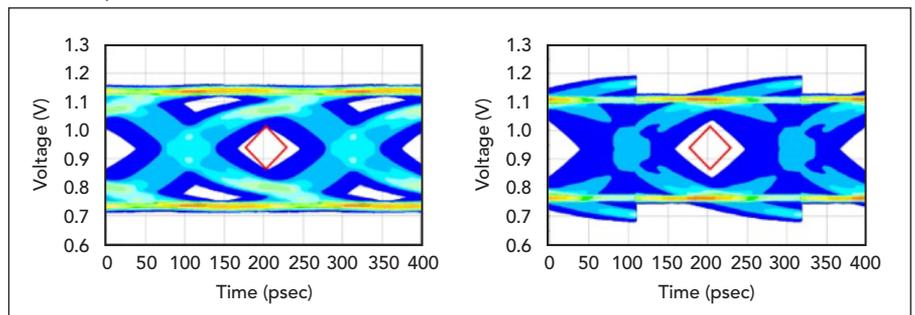
symbol is a digital one, a scaled version of the analog waveform would be added to the original to emphasize the next digital zero. If the detected symbol is a digital zero, a scaled version of the analog waveform would be added to the original to emphasize the next digital one.

Shown on the left of **Figure 3** is an almost closed eye. By applying DFE, the almost closed eye can be opened. As shown in the right of Figure 3, DFE algorithm successfully opens the almost closed eye. Another unique feature of a DFE equalized eye is the kinks before and after the eye opening.

As the data rate increases, one sees a convergence of technologies between serial link and DDR. Before DDR5, no equalization was needed to have a decent eye opening at the receiver. With the push for higher speed and lower power consumption, equalization has become a necessity for an adequate eye opening.

Although it is comforting to have an equalizer at the receiver to improve the eye, one still needs to properly engineer the channel loss and trace impedance so that the equalization can have the most positive impact on the system performance.

To better understand the trade-off between different channel designs and equalization capabilities, the use of electronic design automation (EDA) software as a vir-



▲ **Fig. 3** Left: the eye is almost open with properly engineered channel loss and trace impedance. Right: The decision feedback equalization opens up the eye.

tual prototyping environment has become a necessity as well. By combining the results of virtual prototyping and measurements of the real designs, one forms a robust design workflow that tackles the new and exciting technology.

References

1. JEDEC Solid State Technology Association, DDR5 SDRAM, JESD79-5. Arlington, VA, 2020.
2. Eric Bogatin, Signal Integrity - Simplified. Prentice Hall, 2009.



Accelerate Leading-Edge Memory Design

Streamline your workflow from design to test with Keysight's PathWave ADS High-Speed Digital Design Software.

[Read application note](#)

**Innovators
start here**

Learn more at
www.Keysight.com

Back to Basics: IBIS/IBIS-AMI and the Path to LPDDR5

HeeSoo Lee and Fangyi Rao
Keysight

Starting at the beginning, the core requirement of an SI engineer is to be able to determine whether a data link has sufficient signal integrity. This typically means evaluating the eye-diagram after equalization to see if there is enough margin to achieve a desired bit-error rate (BER). In order to perform this analysis, an engineer needs accurate models of the channel (transmission lines, vias, and other interconnects), and then accurate models of the transmitter and receiver, known as the IO Buffer circuitry and its packaging. However there-in lies a conundrum. Accurate models of the IO Buffer would lead you to the entire SPICE netlist of the IO Buffer, a level of detail that would contain proprietary information about the IC architecture, would contain 1000s of active transistors, and result in very time-consuming simulations.

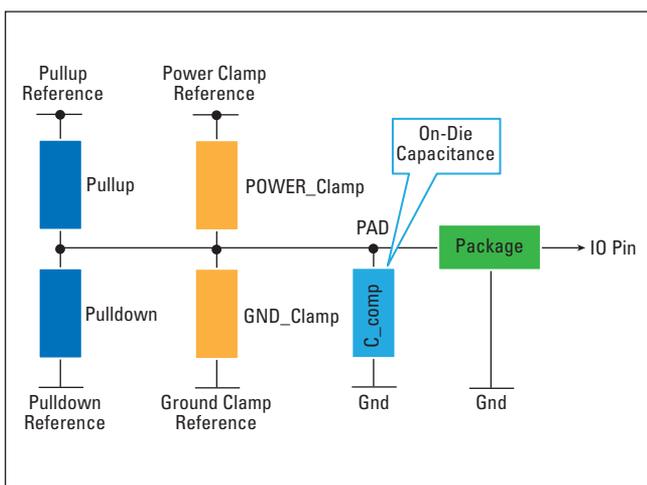
THE BIRTH OF IBIS (I/O BUFFER INFORMATION SPECIFICATION)

The IBIS was released in 1993 to enable silicon vendors, system EDA tools, and simulation end-users to easily exchange models that would protect intellectual property and simulate faster by providing a model that characterizes the analog performance of the IO Buffer into a transportable file. The equivalent block diagram for IBIS is shown in **Figure 1**.

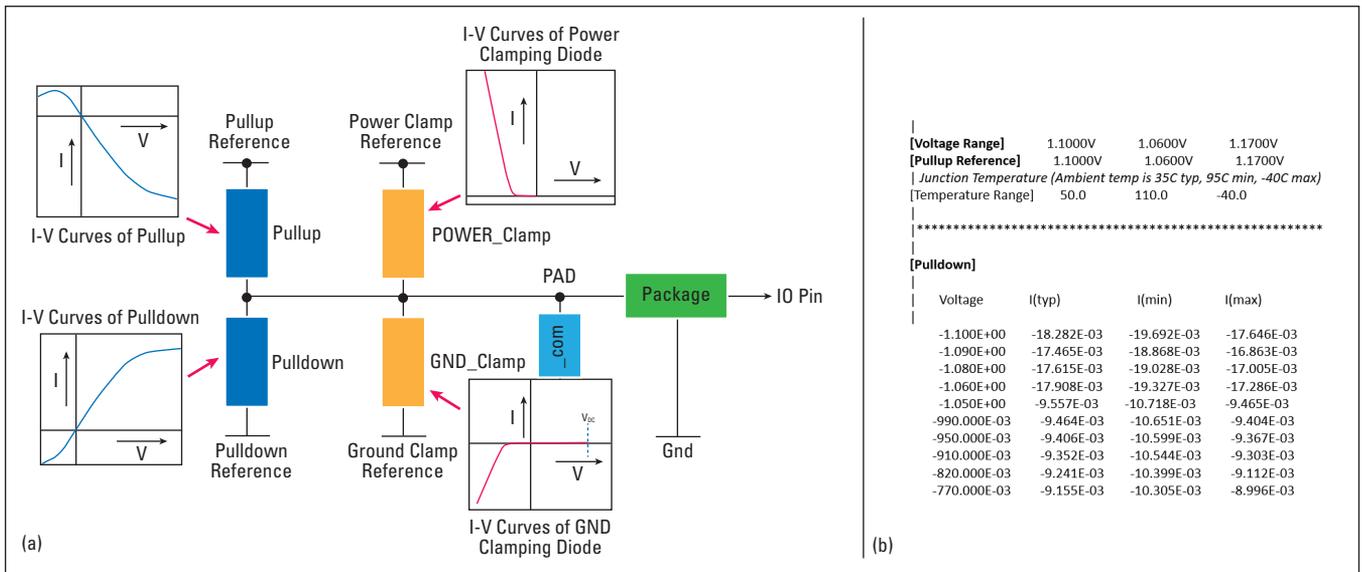
HOW IS THE IBIS MODEL CHARACTERIZED?

An IBIS (.ibs) file is a human readable, text-editable file that contains multiple sets of measured or simulated table-based data representing how the device behaves. In the case of an output model, the data would contain several lists of supply voltage vs. output current (I-V) data for pullup/pulldown and power/GND clamp. This, together with a simply defined 'ramp' slew rate, gives the minimal amount of information a simulator needs. From the I-V tables, the EDA simulator can infer what the current output should be for any channel that we will attach to the output of the IBIS model.

Next, we layer in device behavior for over-voltage and over-current situations. This is done through the Power and Ground Clamp I-V tables, to capture the behavior of the protection diodes found in the IC circuitry. Next, we increase the accuracy of the model with voltage vs. Time (V-T) tables that characterize the exact shape of the rising edges and falling edges as desired (much more detailed information about the waveform than just the slew rate). The V-T tables provide the actual non-linear transition into a known load, which is measured at multiple load conditions.



▲ Fig. 1 IBIS model block diagram.



▲ Fig. 2 I-V table data for IBIS model.

In a nutshell, IBIS models represent I/O buffer behaviors by the table data (I-V, V-T, etc.) from either measurements or simulations, shown in **Figure 2**.

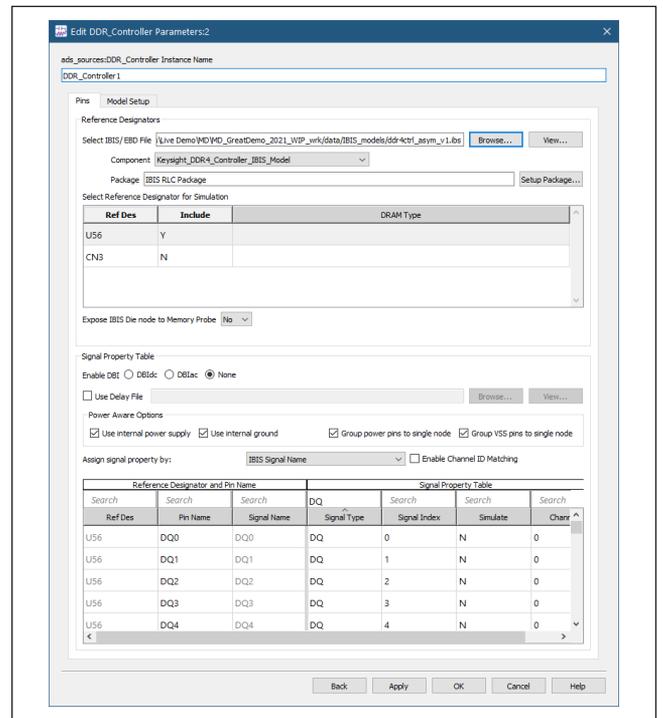
Finally, we can layer in information about the package. At its simplest, this is a description of the typical R, L, and C values for the package pins. It can also be expanded to a definition of R, L, and C for each pin individually, as basic transmission line networks, or as RLC matrices, S-parameters, or SPICE netlists (the latter two in the very latest version 7.0 of the IBIS Specification, ISS – Interconnect SPICE Subcircuit, to capture coupling between pins).

HOW DOES THE IBIS MODEL WORK WITH EDA TOOLS?

So far, that's a lot of information to digest, but luckily, use of a model in the EDA simulator doesn't need an expert knowledge of how the model was created. Entering keywords, data, and making sure the model is compliant to the standards are all the model developer's job. The end-users, consumers of the IBIS model, can easily use the model inside an EDA tool. Typically, users only need to point to the IBIS file, then select the right model for their data rate, the right package model to match their use-case, and the model corner to simulate, shown in **Figure 3**. Corner? - Yes, there is variability in how the IO Buffer silicon would perform from one batch of chips to another. To capture this in the model, IBIS files can contain multiple data sets ('typ,' 'min,' 'max'), for 'typical, fast, slow, min, max' variations, as shown in Figure 2 (b) with an example. An SI engineer is well-advised to run three simulations to check the link performance for typical, fast, and slow model corners to ensure they have enough design margin.

IBIS-AMI (ALGORITHMIC MODELING INTERFACE)

As we have seen so far, IBIS models represent analog electrical behaviors of transmitters and receivers. However, many advanced serializer-deserializer (SERDES) chips employ equalizations such as continuous time linear equalization (CTLE), feed forward equalization,



▲ Fig. 3 Example of IBIS model usage in an EDA tool.

decision feedback equalization (DFE), automatic gain control, along with clock and data recovery (CDR) to compensate the channel loss, inter-symbol interference (ISI), and crosstalk. How does IBIS model handle this?

AMI is the modeling interface for SERDES behavioral models that simulate SERDES functionalities such as equalization and CDR. One example of AMI time domain simulation flow is shown in **Figure 4**. The AMI flow was added alongside the traditional (SPICE-based) IBIS flow in IBIS version 5.0. The AMI portion is specified in a section of the IBIS file known as the [Algorithmic Model] keyword. The combination of the transmitter's analog back-end, the serial channel, and the receiver's analog front-end is assumed to be linear and time-in-

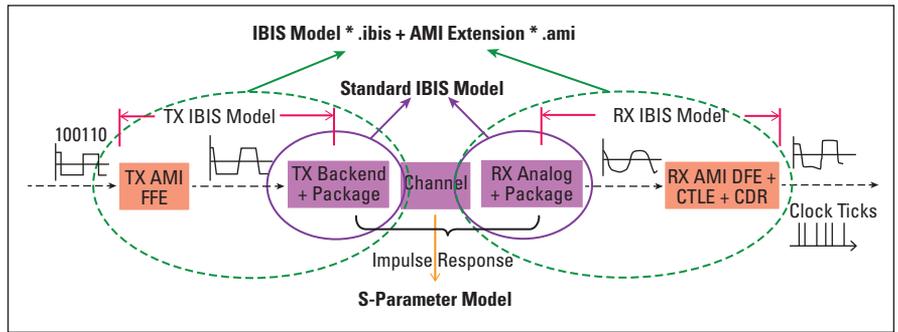
variant. There is no limitation that the equalization should be linear and time invariant in the time domain IBIS-AMI simulation flow. The “analog” portion of the channel is characterized by means of an impulse response leveraging the IBIS constructs for device models. The AMI portion acts as a DSP block which takes an input signal waveform and/or impulse response and outputs a modified waveform and/or impulse response. AMI models are developed by SERDES vendors to match and represent the actual chip behavior. Vendors deliver models in the form of DLL or/and shared object to protect their IP plus the .ami and .ibs text plain files, so that it also provides interoperability between EDA vendors (see **Figure 5**).

Advanced AMI models can perform link training communication to tune the transmitter equalizer parameters for optimized performance and adapt to the signature of any analog channel. This is done when transmitter tap parameters are re-configurable and receivers help them to be configured. Advanced communication specifications such as PCI express, USB, Fibre Channel, and IEEE 802.3 define link training protocols for transmitters and receivers.

If both the transmitter and receiver AMI executable models support the same link training protocol (Back-Channel Interface Protocol), the EDA tool will facilitate the communication between the executable models, enabling link training. Another name for link training in the industry is AutoNegotiation. A link training algorithm can either emulate what the silicon is doing, or it can use channel analysis methods to determine the optimal Tx equalization settings. This ability will also allow Rx AMI models to determine the Tx equalizations settings for channels that do not have automatic link training capabilities.¹

For the model developers, the dynamically loaded executable model implements an application programming interface (API) containing up to five functions: AMI_Resolve, AMI_Resolve_Close, AMI_Init, AMI_GetWave, and AMI_Close. The interface to these functions is designed to support three different phases of the simulation processes: initialization, simulation of a segment of time, and termination of simulation. There are comprehensive programming guides in the IBIS specification.

There are two types of simulations that can be performed with IBIS-AMI models, statistical simulation and time domain simulation, which is also called bit-by-bit simulations. If waveform data is needed for data analysis, time domain simulations must be performed. Traditional spice-like simulations, which are also called transient simulations, can handle complete non-linear behaviors of the system. However, the disadvantage of it is the lengthy simulation time,



▲ Fig. 4 IBIS-AMI time domain simulation flow.

```
[Voltage Range]          1.2000V          1.1400V          1.2600V
[Algorithmic Model]
Executable Windows_c119.00.24215.1_64 DDR_AMI_Rx_x64_v3.d11 DDR_AMI_Rx_v3.ami
[End Algorithmic Model]

[Temperature Range]      60.0             110.0             -10.0

[Add Submodel]
Submodel Name           Mode
34ohm_ODT               Non-Driving

[*****]

[GND Clamp]

Voltage                I (typ)              I (min)              I (max)
-1.20000000E+0         -62.58085700E-3     -70.19394600E-3     -55.60698500E-3
-1.19500000E+0         -61.83645500E-3     -69.45323400E-3     -54.87797600E-3
-1.19000000E+0         -61.09801500E-3     -68.71540400E-3     -54.15423600E-3
-1.18500000E+0         -60.36020700E-3     -67.97844100E-3     -53.43149800E-3
-1.17500000E+0         -58.88904900E-3     -66.50587300E-3     -51.98925300E-3
-1.17000000E+0         -58.15358800E-3     -65.77110700E-3     -51.26990900E-3
-1.16500000E+0         -57.42085500E-3     -65.03634200E-3     -50.55196300E-3
-1.16000000E+0         -56.68872400E-3     -64.30390800E-3     -49.83401700E-3
```

▲ Fig. 5 A usage example of IBIS-AMI algorithmic model.

meaning it's hard to get good, low-level BER calculations. For the IBIS-AMI flow, statistical and bit-by-bit simulations assume the analog portion of IBIS model and the channel to be linear time invariant (LTI). The statistical simulation is based on the impulse response of systems, whereas the bit-by-bit simulation adopts the superposition of single bit responses. With these approaches, the simulations can achieve very low BER calculations at very fast simulation time.

By default, every IBIS-AMI model has an AMI_Init function that allows both the statistical and bit-by-bit simulations. However, in this case, the transmitters and receivers are treated as LTI transmitters and receivers. Therefore, non-LTI features like CDR, gain compression, DFE, clock forwarding, etc. cannot be comprehensively handled with AMI_Init. This is where AMI_GetWave

	Traditional IBIS Flow: SPICE/Transient Simulation	IBIS-AMI Flow: Time Domain (Bit-by-Bit) Simulation	IBIS-AMI Flow: Statistical Simulation
Technology	Nodal Analysis of Kirchoff's Current Laws	Superposition of Single-Bit Response	Calculations Based on Impulse Response
BER Floor in One Minute Simulation	~10 ⁻²	~10 ⁻⁵	~10 ⁻¹⁸ or Lower
Applicability & Assumption	Analog & Channel : NLTV Tx/Rx : NLTV NLTV = Non-Linear and/or Time Varying	Analog & Channel : LTI Tx/Rx : NLTV	Analog & Channel : LTI Tx/Rx : LTI LTI = Linear and Time Invariant

Simulation Speed Setup

Slower ← Harder ← Full

→ Easier → Faster

Non-Linearity Limited

▲ Fig. 6 Summary of IBIS-AMI and IBIS simulations.

function comes in to support those advanced features with IBIS-AMI models. If the GetWave_Exists flag is on, it can handle non-LTI transmitters and receivers. The summary is illustrated in **Figure 6**.

For consumers of IBIS-AMI models, there are four cases or scenarios based on what functions are included in the executable model file. AMI_Init and AMI_Close functions are always in the executable model, meaning that both statistical and bit-by-bit simulations are always applicable. If the non-linear time invariant features are needed, AMI_GetWave must exist and GetWave_Exists flag must be "True" in the IBIS-AMI model, shown in the example in **Figure 7**. (Note that AMI_GetWave only works with time domain or bit-by-bit simulations.)

DDR5 AND LPDDR5 APPLICATIONS

As far as applications for IBIS models are concerned, some of the most complex IBIS models have been created for memory interfaces (DDR). This is due to the large number of signal pins, packages, and configurations available (especially thinking about multiple DRAM dice stacked inside a single package of LPDDR4). Up until DDR4/LPDDR4, IBIS models have covered all the simulation needs of the typical SI engineer.

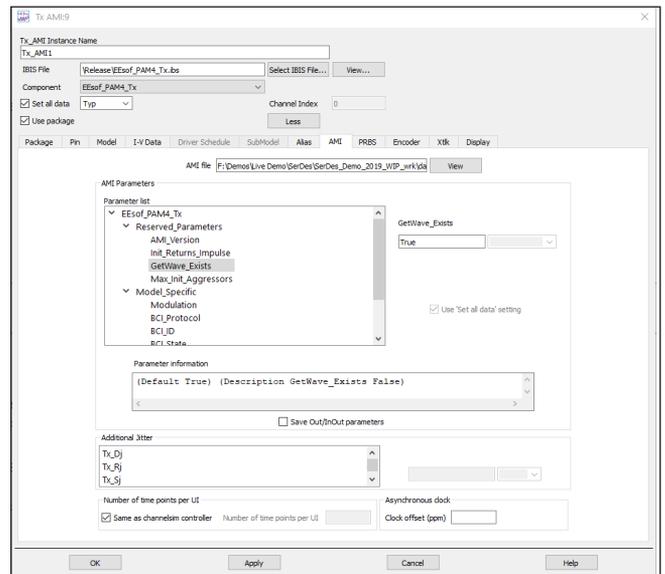
As we move forward to next-generation memories (DDR5/LPDDR5), the technology on chip has evolved, and so must the modeling and simulation technology. In DDR5 and LPDDR5, equalization is available on the commodity DRAM and controller devices for the first time, which came with variable gain, CTLE, and DFE.

The speed in DDR5 and LPDDR5 systems is increased to up to 6,400 MT/s, resulting in worsened ISI impairment. Equalization techniques including deemphasis, CTLE, and DFE are used in memory controller and DRAM to mitigate ISI. Fast speeds also lead to shrunken voltage and timing margins, which are specified at extremely low BER levels. As a result, jitter and noise become critical factors that impact system performances.

In order to produce reliable margin predictions simulations of DDR5 and LPDDR5 systems need to account for the effects of ISI, equalization, jitter, and noise, and millions of bits need to be processed to yield accurate results at specified low BER levels. AMI is a promising candidate as the DDR5/LPDDR5 simulation platform due to its versatility and flexibility in I/O behavioral modeling and its superior simulation speed. However, the unique architecture of DDR channels presents new challenges to AMI when applied to DDR5 and LPDDR5 systems. Recent developments in the AMI methodology have been focusing on addressing these issues, including single-ended signals in DDR channels, asymmetric rise and fall edges in single-ended signals, and clock forwarding.

IBIS-AMI TO SINGLE-ENDED SIGNALS, DDR5/LPDDR5

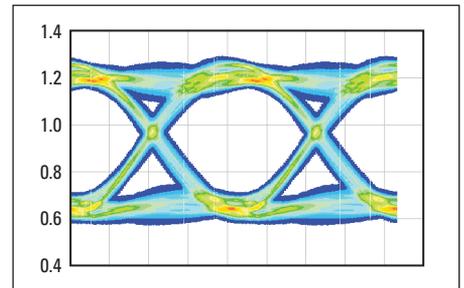
Originally designed for modeling SERDES channels, AMI assumes that all channels are differential and only addresses differential signals. In a DDR channel, data symbols (DQ) and control address command signals are single-ended and have both common and differential components. To resolve this issue, the single-ended input signal to the Rx model is decomposed into a common-



▲ Fig. 7 Example of AMI model with "GetWave_Exists".

and differential component.

The differential component remains the input waveform to the Rx AMI_GetWave function, which is the same as in the current



▲ Fig. 8 Asymmetric rise and fall edge of DDR signals.

The common component, which is assumed to be a constant, is characterized by the EDA tool as the mean value of the steady state high and low voltages at the Rx pad. The value is passed to the Rx model by the EDA tool in the AMI_Init call through a new DC_Offset parameter. In the AMI_GetWave function the Rx model can choose to internally recover the single-ended input signal by adding DC_Offset to the differential input waveform.

ASYMMETRIC RISING AND FALLING EDGES OF SINGLE-ENDED DDR SIGNALS

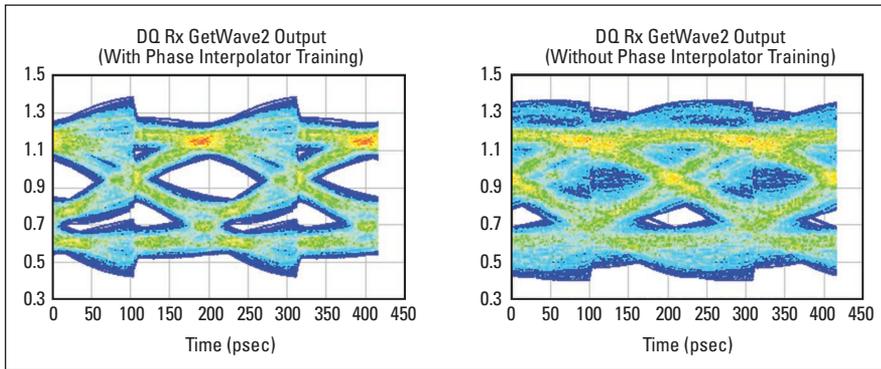
AMI also assumes that rise and fall edges are symmetrical in the signal. While this may be a valid assumption for differential I/O, it is typically not the case for single-ended I/O, where the pullup and pulldown slew rates are usually noticeably different. As a result of asymmetric edges, the single-ended eye is asymmetrical vertically, and its crossing level is shifted either upward or downward from the center voltage of the eye, impacting both voltage and timing margins. To capture these effects, advanced AMI simulation algorithms are developed to take into account the difference between rise and fall waveforms.

Figure 8 shows a DQ eye at the Rx pad generated by an AMI simulation. In the plot, the rise and fall edges are asymmetric as is typical for a single-ended signal, and the crossing level is shifted upward from the center voltage of the eye due to the asymmetric nature. Note that

Figure 8 also shows the DC offset of the single-ended DQ signal.

NEW FORWARDED CLOCKING SOLUTION

In the AMI specification, it is assumed that every Rx has its own CDR circuitry to recover the clock from the data, and the AMI_GetWave function has only one input waveform, which is the data signal. However, DDR channels employ the so-called clock forwarding architecture, where, instead of using an internal CDR, the DQ Rx uses a data strobe signal (DQS) as the forwarded clock to clock the DQ Rx DFE slicer and data sampling. Practically, the DQ Rx device has two input signals, one is data, and the other is clock. To enable modeling of clock forwarding, a new Rx AMI_GetWave API, originally known as GetWave2, is established in IBIS BIRD 204 and approved for a future release of IBIS specification. The API defines two input waveforms for data and clock signals, respectively. The DQ Rx clocking behavior can be physically modeled in the new AMI_GetWave function.



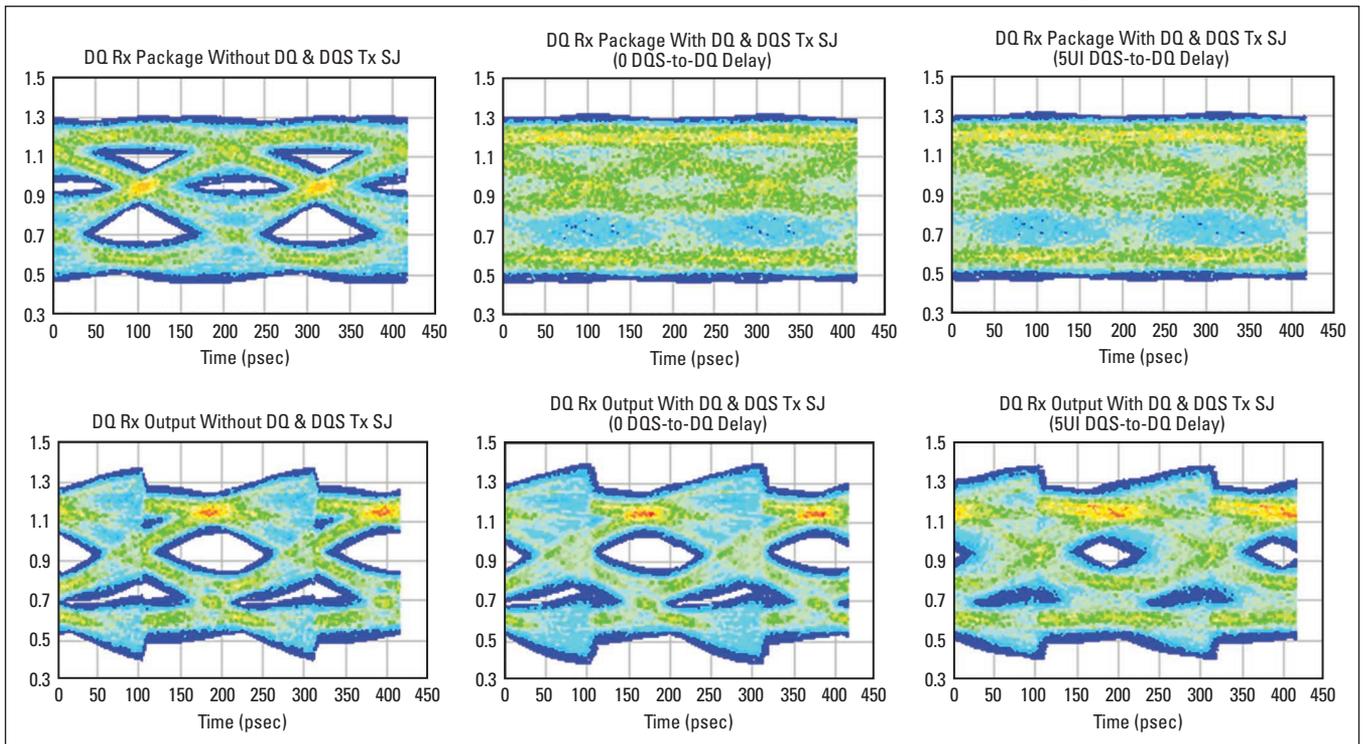
▲ Fig. 9 Before and after phase interpolator training with forwarded clock modeling.

PHASE INTERPOLATOR IN FORWARDED CLOCKING

Besides clock forwarding, another key clocking functionality that can be modeled using the new AMI_GetWave API is the phase interpolator in the controller DQ Rx. During READ cycles, the controller DQ Rx PI applies a 90-degree phase shift to the forwarded DQS signal and mixes it with the original one. The resulting signal is a delayed DQS signal, and the delay value depends on the mixing weights. During system training, the controller tunes the weights and, therefore, the delay to adjust the DQ-DQS skew for optimal DQ Rx DFE clocking in READ operations. **Figure 9** shows a READ cycle controller DQ post-DFE eye with and without PI training modeled by the new AMI_GetWave API. The training aligns DFE switching with data bit edges to help open the eye.

JITTER TRACKING WITH FORWARDED CLOCKING

One advantage of the clock forwarding architecture is jitter tracking. Because the DQS signal is used to clock the DQ Rx, when the DQ is sampled, correlated jitter between DQ and DQS are cancelled. On the other hand, the DDR5 spec allows a certain amount of electrical path mis-match between DQ and DQS Rx. The mis-match reduces the DQ-DQS jitter correlation and adversely impacts the effectiveness of jitter tracking and DFE. With the new AMI_GetWave API, both jitter tracking and the effect of unmatched Rx can be captured naturally in AMI simulations. **Figure 10** shows simulated eyes of a DQ signal at the Rx package pin and at the Rx DFE output.



▲ Fig. 10 Jitter tracking.

Without Tx jitter, the eye is almost closed by ISI at the package but opened by the DFE at the Rx output. When SJ is injected to DQ and DQS Tx, the eye is completely closed at the package. In the case of matched Rx (with zero DQS-to-DQ delay) DQ and DQS jitters are correlated and tracked by DQ sampling times, leaving the DQ post-DFE eye almost unchanged from that without Tx SJ. In the case of unmatched Rx (with a 5UI DQS-to-DQ delay) the DQ-DQS jitter correlation is reduced, and the jitter tracking becomes less effective, leading to a worsen DQ post-DFE eye.

CONCLUSION

In this article, we reviewed the basics of IBIS and IBIS-AMI models. IBIS/IBIS-AMI models are very effective vehicles for chip vendors to communicate and share their intellectual property with customers without harming their design secrets. Also, from the system vendor's point of view, it is the fastest and easiest way to evalu-

ate and validate their designs instead of going through multiple board spins. That is why IBIS/IBIS-AMI models have been very popular in high-speed digital designs and became the market standard for DDR and SERDES applications.

Due to the ever increasing speed-grade of memory systems, it is necessary to apply equalizations, which creates severe burdens for memory system design engineers. Fortunately, the challenges have been overcome by an IBIS-AMI solution for single-ended signals and the introduction of a forwarded clocking solution in BIRD 204. We anticipate new challenges when the next generation of memory systems comes, such as (LP)DDR6 or GDDR7, but we can count on new solutions coming out to help design engineers.■

References

1. IBIS specification version 7.0.



Get Ready for DDR5 and the Road Ahead

Discover how to anticipate emerging test challenges and succeed in next-generation memory design.

Get the white paper

Innovators
▶ start here

Learn more at
www.Keysight.com

Managing PCB Crosstalk

Donald Telian
SI Consultant, SiGuys

“Crosstalk” occurs when energy in one signal (called an “aggressor”) couples onto another signal (the “victim”), adversely affecting the victim signal’s performance. The aggressor/victim language associated with crosstalk indicates danger is lurking, provoking hardware engineers to constant vigilance. How can we tame this foe? Or, more specifically, what causes crosstalk? When does it become problematic? What can you do to ensure it does not ruin your product design? I’ll answer those questions in a moment, but first let’s have a look at the crosstalk issue I most often find and correct in today’s designs.

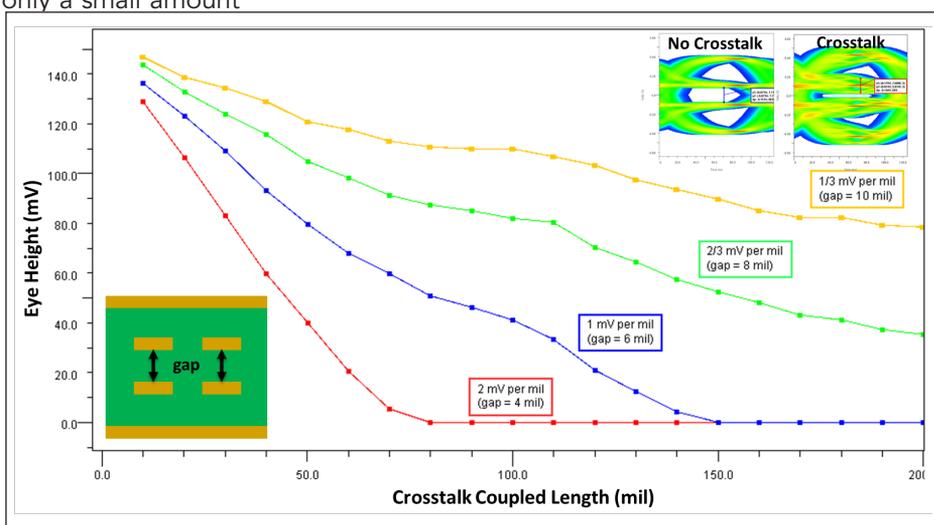
THE MOST COMMON CROSSTALK ISSUE

As design tools and practices have matured, the most common crosstalk issue that escapes a design team’s notice is vertical layer-to-layer coupling. While solid planes are used to prevent this, voiding in these planes places small holes that signals can couple through. In my experience, these “Z direction” couplings are not found by design rule checks, and it requires only a small amount of vertical coupling to collapse and eye. This issue is growing because connector and capacitor pads are increasingly of a relevant feature size [RFS, 1] that must be impedance matched. More holes, more chances for coupling.

Figure 1 quantifies the impact of vertical layer-to-layer crosstalk on PCIe Gen3 eye height when a ground shield layer is not in place. The eye diagrams in Figure 1 show performance without crosstalk (left) and with crosstalk (right). Because the link is short (3”), the signals are over-equalized, and hence four voltage levels are seen (yes, this happens quite often [2]). With no crosstalk, the eye opening is ample at 150 mV. With crosstalk,

each of the four voltage levels are widened by ~150 mV of noise, closing the eye. Simulated eye heights are plotted in the graph, revealing how eye height decreases as the amount of coupled parallelism between the layers increases from 0 to 200 mils on the X axis. As the “gap” distance between the layers decreases (gold=10 mils to red=4 mils, in 2 mil increments), eye height decreases at the rates shown in the color-coded boxes. These curves are easily created in Signal Integrity Toolbox™, so try out this exercise on your design using a [free trial of the software](#).

Figure 1 reveals inter-layer crosstalk can cause a 1 mV decrease in eye height per 1 mil of coupling when the layer-to-layer gap is 6 mils (blue). That means only 100 mils of parallelism can remove a generous eye margin. So, make sure diff-pairs do not overlap through gaps in ground planes – which typically must be done as a manual/visual process. That said, let’s take a step back and explain both the sources of crosstalk and the design methods that prevent it.



▲ Fig. 1 Inter-Layer Crosstalk on PCIe Gen3 Eye Height, versus Layer Gap and Coupled Length (plots created in MATLAB and [Signal Integrity Toolbox](#))

TABLE 1.

FACTORS CONTRIBUTING TO CROSSTALK, THEN AND NOW.

Factor	1980	2020	Notes/Relevance
Signal Spacing	15 Mils	4 Mils	How close is victim to aggressor?
Parallel Length to Saturate	3 Inch	Any	Increases with parallelism, to saturation.
Voltage Swing	5 V	1 V	How much voltage can couple?
Rise Time	10 ns	20 ps	3 orders of magnitude in di/dt and dv/dt.
Distance to Ground	Hasn't Changed		Will signal couple to victim or ground?
Typical Route Lengths	Hasn't Changed		How long can coupling occur?
Typical Crosstalk	2%	30%	Without managing a design's crosstalk.

THE MECHANICS OF CROSSTALK

Over the years technology has worked against us causing typical (unmanaged) crosstalk voltages to increase from 2% to 30%, as shown in **Table 1**. As data rates increase and voltage margins decrease, even the smallest, unexpected signal disturbance becomes problematic – even just a few millivolts. As such, it's important for engineers working in all aspects of electronics design and production to have a basic understanding of the mechanics of crosstalk.

Table 1 lists the factors that contribute to crosstalk. Intuitively, the closer the signals are to each other the greater their potential for coupling or crosstalk. As signals travel “close” together over increasing length (referred to as “parallelism”), the amount of crosstalk increases to the point of “saturation;” at saturation the maximum amount of crosstalk has been reached. As shown in Table 1, modern technology saturates very quickly so we don't think about this as much as we used to. Crosstalk also grows with both voltage swing and rise time, or with increasing dv/dt and di/dt. In terms of the familiar equations, $I=C*dv/dt$ and $V=L*di/dt$, capacitance increases as metal moves closer together and so does mutual inductance – and hence all factors continue to combine and increase crosstalk. As such, controlling signal spacing (and, if possible, voltage swing and edge rate) directly impacts the magnitude of crosstalk in your design.

To understand how the factors interact and which factors are dominant, try entering the values in Table 1 into this [on-line crosstalk calculator](#) [3] ($H=10$ mils, $h1=h2$). Modify the parameters and observe what changes – this will enhance your crosstalk intuition. Perhaps try out the values inherent in your design.

Surprisingly, despite this increase in crosstalk potential, we have seen an overall decrease in issues. How can that be? Like other design challenges, the technology world rallied with awareness of the problem, de-

signed rules to prevent it, and designed tools to ensure those rules are followed. So, before we panic, let's put crosstalk problems into perspective.

CROSSTALK IN PERSPECTIVE

Yes, crosstalk problems are real, but you might be surprised to learn I've encountered only three serious issues in 40 years – designing all types of electronic products. All three issues were found after hardware was built and fueled new disciplines in preventing crosstalk problems prior to implementation. As the issues are instructive, let's take a look at what caused them.

As stated previously, the leading cause of system-level crosstalk failures is unshielded layer-to-layer parallelism in the Z (vertical) direction. Indeed, this caused two of the three problems. One was a long section of parallelism between a “high-speed” signal and a “low-speed” signal (watch out for this, “low-speed” signals don't get enough attention anymore). The other problem involved two serial links signals with only 100 mils of coupling through plane cutouts. Both problems were extremely difficult to isolate, with the “aha” moments occurring during a careful study of layer-to-layer PCB layout artwork. While layout tools may assert they DRC (design rule check) these situations, I still visually overlay and examine adjacent layers for potential issues – particularly around cutouts. This is a situation where brain-power and experience surpass the capabilities of computer algorithms.

The third crosstalk issue was in package-level bond wires caused by interleaved inputs and outputs buffered within the IC. Crosstalk induced the inverse of the output back onto the input, and the resulting oscillation was so powerful and predictable I applied for a patent on this novel oscillator design. Who says problems can't become inventions?

Because crosstalk problems are difficult to isolate and correct in hardware, and hence severely impact a

product's performance and schedule, majority of engineers simply design it out – albeit with increasing material cost. The exception to this might be very high-volume products; these design teams use detailed simulations and manual layout to minimize cost. But again, most product implementation teams simplify and solve the crosstalk problem by using design rules.

CROSSTALK DESIGN RULES

Crosstalk design rules reduce crosstalk to acceptable levels by managing the two directions in which signals can couple within a PCB: vertical and horizontal. Vertical crosstalk is caused by signals on other layers, or "inter-layer." Horizontal crosstalk is caused by signals on the same layer, or "intra-layer." Crosstalk from each direction is handled in different ways, as follows:

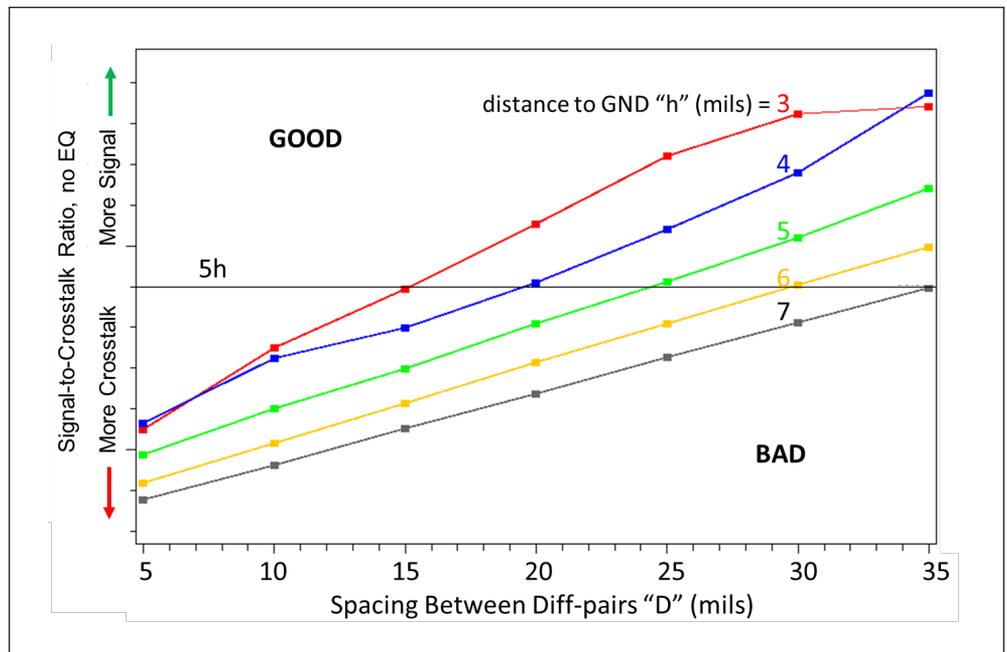
Inter-layer Crosstalk

Inter-layer crosstalk problems are prevented by placing solid ground planes (shields) between signal layers. Although adding layers adds cost, solid planes solve numerous SI problems such as controlling trace impedance, return current, power supply impedance, and bypass capacitor loop current. So extra ground layers are readily added in all but the highest volume products. This sounds simple enough but be advised that a "solid" plane never exists in practice. As such, I'll stress again it's important to verify that signals will not couple through cutouts, antipads, or other gaps in the plane. In these areas, signals on both sides of the "shield" remain susceptible to crosstalk because part of the shield has been removed.

Intra-layer Crosstalk

Intra-layer crosstalk is prevented by enforcing a spacing distance between signals greater than $5h$ to $7h$, where " h " is the distance between the signals and their adjacent ground plane(s). The design rule is stated in terms of " h " to ensure the signal's coupling to a nearby plane (which is good) is roughly an order of magnitude greater than its coupling to a nearby signal (which is bad). In practice, this generally requires signals to be spaced about 25 mils apart.

To illustrate the efficacy of the " $5h$ " design rule, **Figure 2** shows a crosstalk signal-to-noise ratio on the Y axis versus the spacing distance " D " between two signals on the X axis. As the Y axis is a ratio (not detailed here), larger values are "good" and smaller values are "bad" as shown. The colors show " h " (the stripline



▲ Fig. 2 Intra-Layer Crosstalk Magnitude versus Signal Spacing and Distance to Ground.

trace's distance to ground in each direction) varying from 3 mils (red) to 7 mils (black), in 1 mil increments. The horizontal line marks a constant magnitude, which is the $D=5h$ location for all values of h . For example, the $h=3$ mil line (red) crosses the horizontal line at 15 mils, the $h=4$ mil line (blue) crosses at 20 mils, and so on. While minor non-linearity is seen with small h values, the plot demonstrates how the design rule achieves a consistent crosstalk ratio across a variety of stackups and implementations.

Figure 2 illustrates both how signal quality increases (i.e., decreasing crosstalk) as spacing between signals increases (larger D), and how an acceptable crosstalk level can be reached sooner if signals are closer to ground (smaller h). Again, manipulating " D " and " h " is the primary mechanism for controlling intra-layer crosstalk. Consult the design guidelines associated with your components or technology to determine the recommended D/h ratio; I expect you'll find it to be in the in the 5 to 7 range, unless a constant D is used instead.

As stated previously, automated layout tools are better at enforcing intra-layer than inter-layer spacing rules. As such, ground shields are typically used vertically, and spacing rules are used horizontally. In rare situations, ground moats have been used horizontally and spacing rules are used vertically, and the physics involved is similar to that described above.

While it's best to prevent problems before they happen, when confronted with crosstalk in hardware don't forget you likely have programmatic control over SerDes/DDRx drive strength, edge rate, and equalization. You may find you can fix the problem using software [2]. For example, simply turning off the Tx equalization shown in Figure 1 can restore the eye – even without removing the crosstalk.

CONCLUSION

Crosstalk problems can be real yet are not necessarily as pervasive as one might expect as long as design best practices are followed. Here we've discussed the factors that exacerbate crosstalk and how to manage them using design rules. Crosstalk simulation is used to develop physical design rules that are simple to implement, and also to cross-check and adapt the rules for a specific PCB when lowest cost is desired.

This article is an excerpt from Donald Telian's book "[Signal Integrity, In Practice](#)." A Practical Handbook for Hardware, SI, FPGA, and Layout Engineers.

References

1. Telian D. (2022 April 1). '[Which Discontinuities are Small Enough to Ignore?](#)' Signal Integrity Journal RSS.
2. Telian D. (2022 May 3). '[Fixing Signal Integrity Issues in Software](#).' Signal Integrity Journal RSS.
3. (2022 September 29). "[Stripline Crosstalk Calculator](#)." EEWeb's PCB Tools.

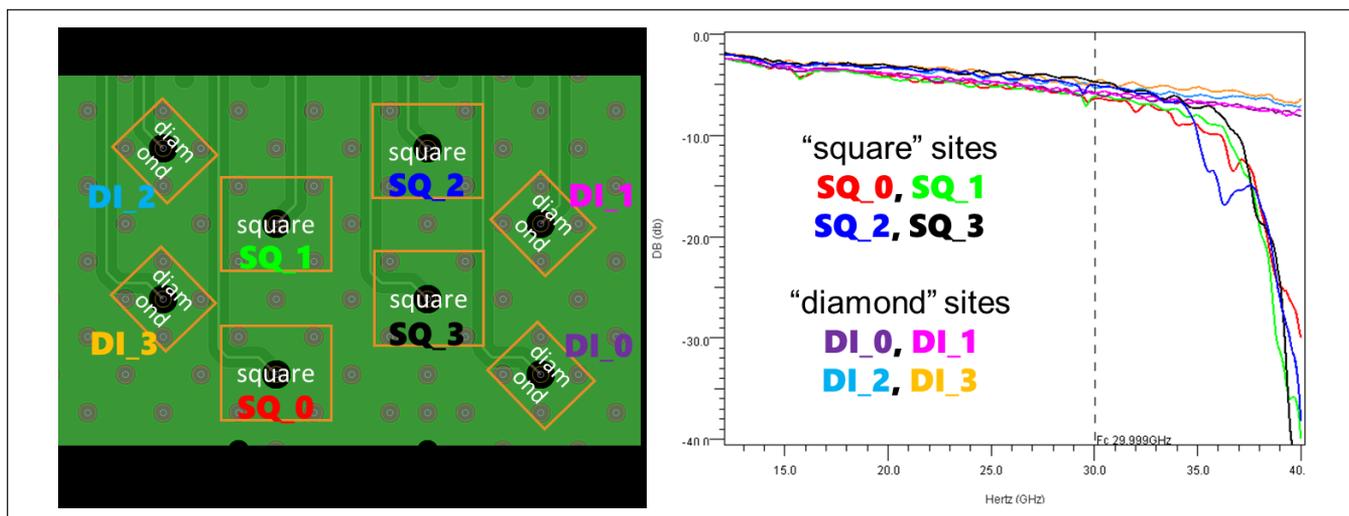
Proper Ground Return Via Placement for 40+ Gbps Signaling

Michael Steinberger, Donald Telian, Tsuk Michael, Iyer Vishwanath, and Yanamadala Janakinadh
MathWorks and SiGuys

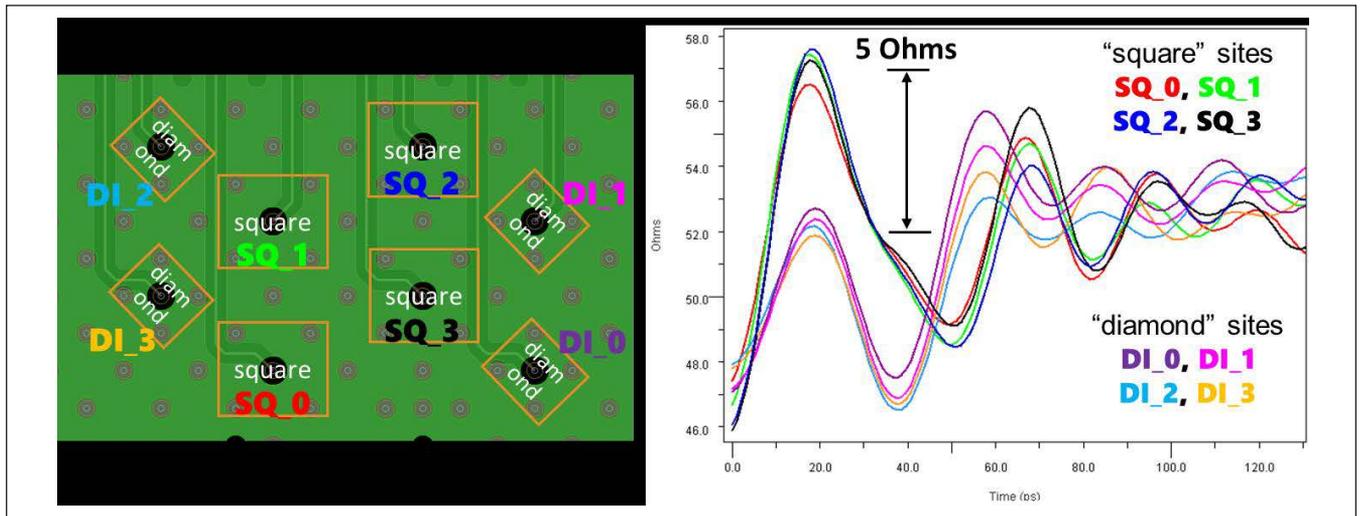
While the physical design and manufacture of electronic systems has advanced significantly over the years, changes in dimensions and density in printed circuit boards (PCBs) have been incremental – particularly compared to the exponential increase in integrated circuit (IC) density and system interconnect data rates. Indeed, in the past 30 years IC density has increased 100,000x while PCB density has increased 3x [1, page19]. As such, a challenging convergence of operating frequencies and standard PCB dimensions looms on the horizon. For example, although effort is made to keep 28+ Gbps via stub lengths less than 5 mils, few recognize the surface mount pad stub extending beyond the backside of soldered connector

pins is often significantly longer than 5 mils.

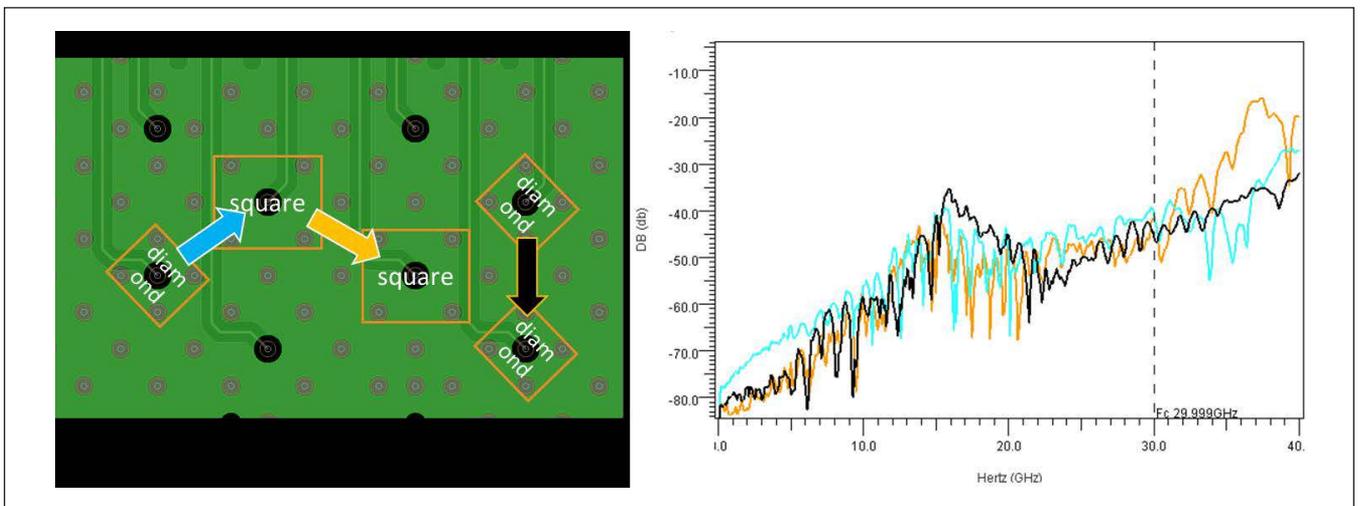
Of particular concern and focus in [this paper](#) is the placement of ground return vias (GRVs) near signal vias. For decades hardware and layout engineers have added GRVs near signal via layer transitions based on best practices, folklore, and fear with little understanding of where GRVs need to be and why. Mystery and misunderstanding have led to re-routes and wasted PCB real estate. As data rates continue to increase, driving significant spectral content into the 40GHz to 60GHz region, it will become ever more important and difficult to place these GRVs where they will get the job done. The goal of [this paper](#) is to describe the role and behavior of GRVs in a way that informs design and layout engineers’



▲ Fig. 1 PCB test structure and resulting signal IL.



▲ Fig. 2 Measured TDR for eight test signal vias.



▲ Fig. 3 Measured crosstalk for signal vias.

intuition and engineering judgement, using practical examples.

Figure 1 shows eight single ended signals under a 1 mm pitch ball-grid array with equivalent ~ 100 mil vias to the route layer shown, next to measured data for the same. Each of the eight signal vias is immediately surrounded by four GRVs. However the pattern of the GRVs varies, depending on where each signal falls within an alternating 2 mm array of GRVs. The GRVs in the sites labeled "diamond" are closer to their associated signal via (1 mm) than the GRVs in the sites labeled "square" are to their associated signal via (1.4 mm). Measured data (at right) reveals the "square" signal's insertion loss (IL) decreases to -40 dB at 40 GHz while the "diamond" signal's IL continues to decrease linearly. How is it possible, simply due to ground via placement, that 99% of the signal is lost for half of the signals while IL for the other half is well-behaved? Furthermore, how can a tiny 100-mil via structure within the same dielectric material exhibit more loss than fifteen inches of trace? [This paper](#) will demonstrate the answers lie in understanding the interactions of the via's return currents.

[This paper](#) will demonstrate what happens when GRVs are too distant for the data rate at hand. When the distance from the signal via to the GRVs in Figure 1 becomes greater than approximately a quarter wavelength, the structure resonates with a relatively high Q. It is in effect a microwave filter. To help avoid the excessive IL shown in Figure 1, [this paper](#) will define a Gap-Rate Distance (GRD) metric that can be easily applied to GRV placement in a PCB layout.

We will use three metrics in both simulation and measurement to gauge the efficacy of passive interconnects as influenced by GRV placement: IL, time domain reflectometry (TDR), and crosstalk. While IL currently gets the most attention because it both reduces the signal amplitude and is a major source of intersymbol interference, SerDes equalization schemes and lower-loss materials have been effective at mitigating IL effects. However, as the authors have been asserting for a long time^{2,3,4}, transmission path discontinuities, as measured by TDR, are every bit as serious a source of intersymbol interference, and significantly more difficult to equalize. Indeed, as increasing miniaturization impacts electron-

ic products, discontinuities are becoming the primary cause of link failure.^{1, Chapter 4}

Measured TDR in **Figure 2** demonstrate that the GRV placements that affect IL in Figure 1 also cause unexpected discontinuities in the transmission path. For ~15 ps, a perturbation relevant to current data rates, the signal via impedance is consistently five ohms higher for the “square” sites than the “diamond” sites. Although one such increase in via impedance might not be a serious problem, multiple irregular vias along a transmission path can cause serious impairments.

Finally, as the measured data and model results in [this paper](#) will show, crosstalk is going to become a serious impairment as data rates increase. As shown in **Figure 3**, measurements demonstrate that at higher frequencies crosstalk between the “square” sites (gold) is higher than the crosstalk between a “square” site and a “diamond” site (blue), which in turn is higher than crosstalk between two “diamond” sites (black). Crosstalk increas-

es rapidly with frequency and is primarily a function of GRV configuration. Note also from the layout that the signal vias are not at all “close” to each other compared to crosstalk dimensions normally considered, indicating some effect beyond capacitive coupling is at work, as will be demonstrated. Although differential transmission improves the situation somewhat, a similar phenomenon occurs for that case as well.

[This paper](#) is a natural extension of the authors’ compute-efficient and structure-based approach to via modeling^{5,6,8} to include effects of higher frequencies. While many current applications use differential transmission, [this paper](#) will concentrate on single-ended transmission because the role of the GRV is simpler to illustrate and comprehend. However, we will briefly address differential transmission in section 7.1 near the end of the paper.

The paper referenced here received the Best Paper Award at [DesignCon 2022](#). To read the entire DesignCon 2022 paper, [download the PDF](#).

Keysight enables innovators to push the boundaries of engineering by quickly solving design, emulation, and test challenges to create the best product experiences. Start your innovation journey at www.keysight.com.

This information is subject to change without notice. Keysight Technologies, 2023, Published in USA, August 3, 2023, 7123-1074.EN



Maximize Your PCB Design Productivity

Quickly identify signal integrity issues in your PCB layout with Keysight’s EP-Scan Software.



Innovators
▶ **start here**

Learn more at www.Keysight.com

[Read application note](#)