gEEk® spEEk

Successful PCIe® Interconnect Guidelines at 8, 16, and 32 GT/s

Presenter: Steve Krooswyk

Introduction

- We face an increasing pressure to improve system interconnect at each data rate bump
- Optimizations for loss, reflection and crosstalk are up for consideration
- But we have questions, how much is enough?
- And when should we run channel simulation?
- Let’s answer these questions and review high speed design considerations that apply to PCIe®

Agenda

- High Speed Interconnect Reminders: Transmission Lines & Vias
- PCIe® Connector and Cable Considerations
- Solution Space Discussion: Paper Calculations or Simulation?
- Simulation Examples and Topology Solutions

PCIe® Generations

	Data Rate	Total Budget	Add in Card Budget	Reach Goal
PCIe® 3.0	8 GT/s	23.5 dB	6.5 dB	20-inch and standard PCB
PCIe® 4.0	16 GT/s	28.0 dB	8.0 dB	16-inch on mid-range PCB material
PCIe® 5.0	32 GT/s	36.0 dB	9.5 dB	14-inch on low-loss PCB material

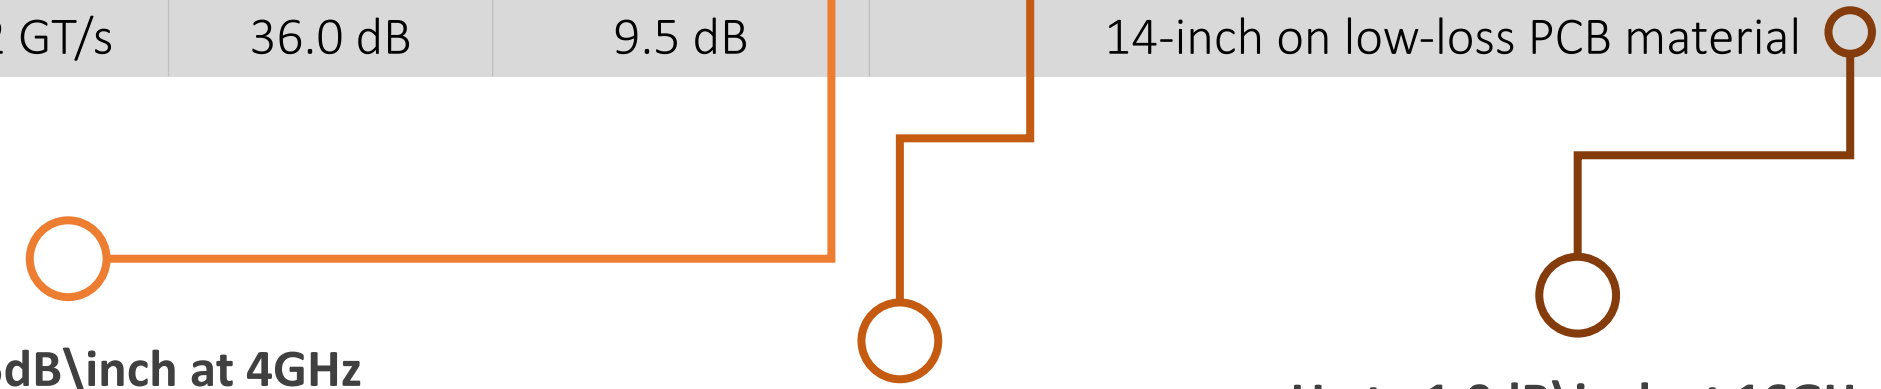

Up to 0.85dB\inch at 4GHz
Dissipation factor > 0.018
Standard FR4

Up to 1.4dB\inch at 8GHz
Dissipation factor < 0.015
IT180
370HR
etc.

Up to 1.0dB\inch at 16GHz
Dissipation factor < 0.010
Megtron 6
I-Tera MT40

Transmission Lines

Glass Weave Management

- Skew appears on SDD12 at 2 GHz
- Zigzag routes or route @ glass pitch

gEEk® spEEk May 21, 2020 – Quantifying Glass Induced Skew on PCBs

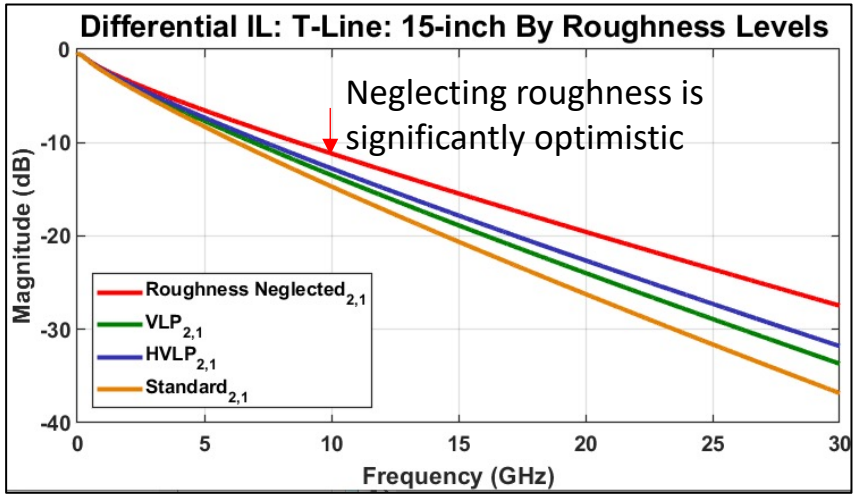

Surface Roughness

- Smoother copper may not be required on shorter, lower speeds
- Significant impact at Gen4+
- Evolving modeling methods (DesignCon papers)

15" Transmission Line on mid-range material: ER=3.3, tanD=0.010

	PCIe® 3.0 - 4 GHz	PCIe® 4.0 - 8 GHz	PCIe® 5.0 - 16 GHz
Standard Foil (7-14um)	-6.9 dB	-12.3 dB	-21.8 dB
VLP (1-3um)	-6.4 dB +7%	-11.2 dB +9%	-19.9 dB +9%
HVLP (<1um)	-6.1 dB +12%	-10.6 dB +14%	-18.8 dB +14%

Transmission Lines: The Bottom Line on 85 Ohms

gEEk spEEk

- 85 ohms required for interoperable systems and cards
- 85 ohms does not apply to packages
- PCB for captive, or custom, systems do not require 85 ohms
- Connectors do not require 85 ohms
- Cables are not bound to 85 ohms

4.7.8. Differential Data Trace Impedance

The PCB trace pair differential impedance for a 5.0 GT/s capable data pair must be in the range of 68 Ω to 105 Ω. The PCB trace pair differential impedance for an 8.0 GT/s capable data pair must be in the range of 70 Ω to 100 Ω. These limits apply to both the add-in card and the system board.

Notes
- Motherboards with long (high loss) channels may need to have tighter impedance control.
- This requirement does not apply to vias, the connectors, package traces, cables, and other similar structures.
- Designs should still attempt to minimize the impedance discontinuities from vias, the connectors, package traces, cables, and other similar structures.

IMPLEMENTATION NOTE

Differential PCB Trace Impedance

The PCB trace impedance requirement specified in Section 4.7.8 only applies to topologies that support 5.0 GT/s or 8.0 GT/s covered by this form factor specification that use the connector defined in this form factor specification.

Specifically, the *PCI Express Card Electromechanical Specification* covers the following two topologies (as defined in Section 4.6.1):

PCI Express devices across one card electromechanical connector on a system with a system board and an add-in card

PCI Express devices across two card electromechanical connectors on a system with a system board, a riser card, and an add-in card, where the connector between the riser card and the add-in card is a card electromechanical connector.

Other topologies governed by different specifications may impose different impedance requirements or leave the impedance unspecified.

For example, the topology of "PCI Express devices on the same system board" does not fit within a form factor specification and hence must only follow the requirements of the *PCI Express Base Specification*. The *PCI Express Base Specification* does not define a PCB trace impedance requirement so with this topology designers can choose the PCB trace impedance that is best for their applications.

Add In Card (CEM)

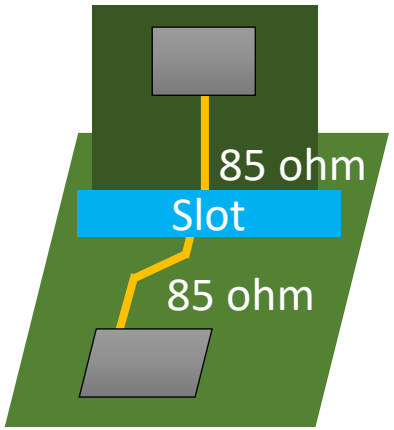

Captive Link
(Everything Else)

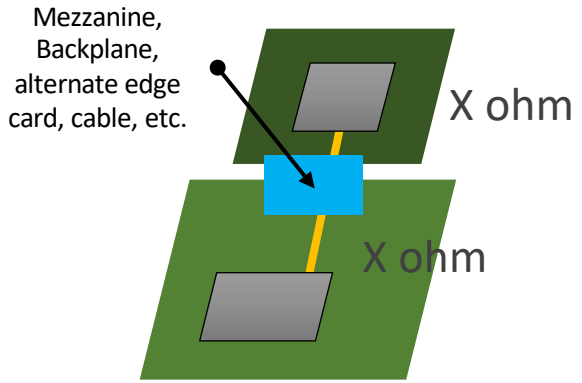

Transmission Lines

- Solder mask: conformal vs. blanket
  - 1 ohm impedance
  - A brief model review has found impact loss by 2%, crosstalk by 4%
- Plating height – microstrip is not all trapezoidal

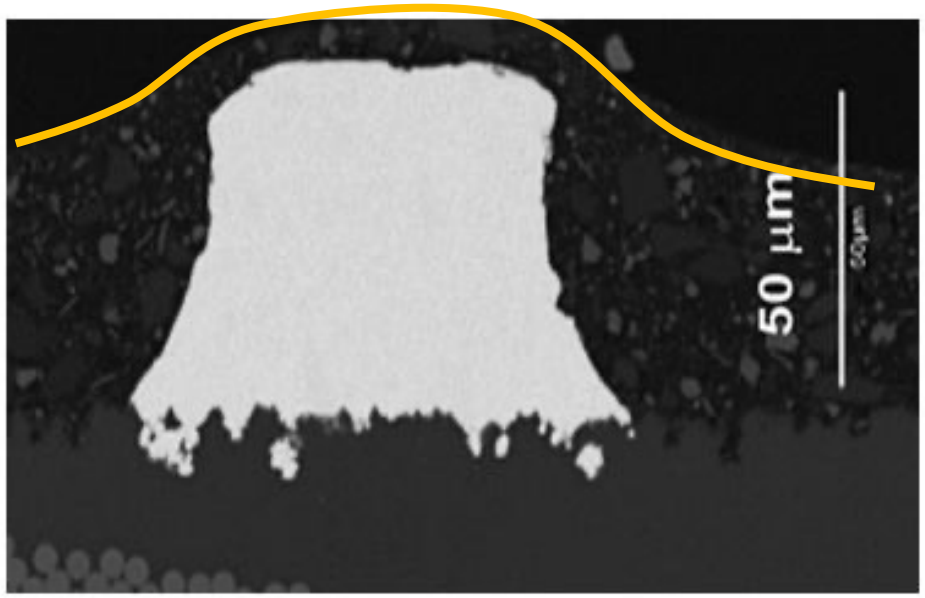

Vs.

Transmission Lines: Modeling Metal Layer

- Metal layer fills with resin from prepreg during press
  - Expect dk for pockets between 2.8-3.0
- Prepreg properties on metal later is optimistic loss and crosstalk
- TDR: 1-2 ohms difference

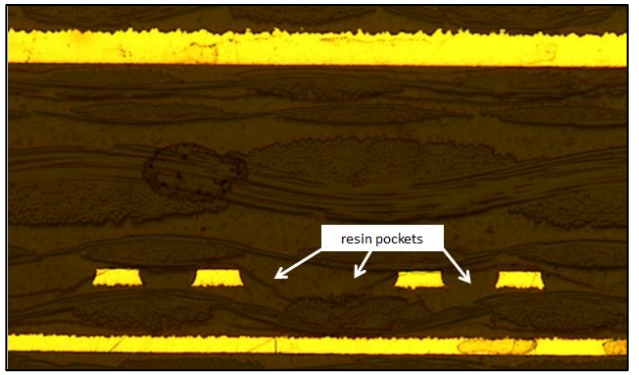

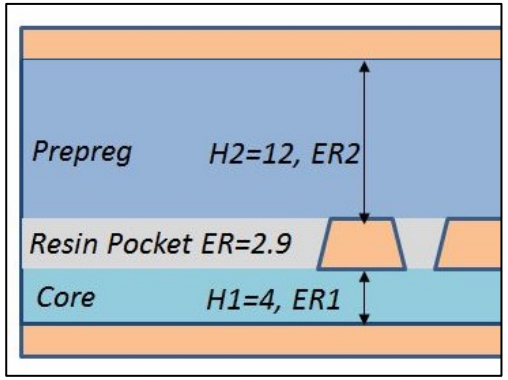

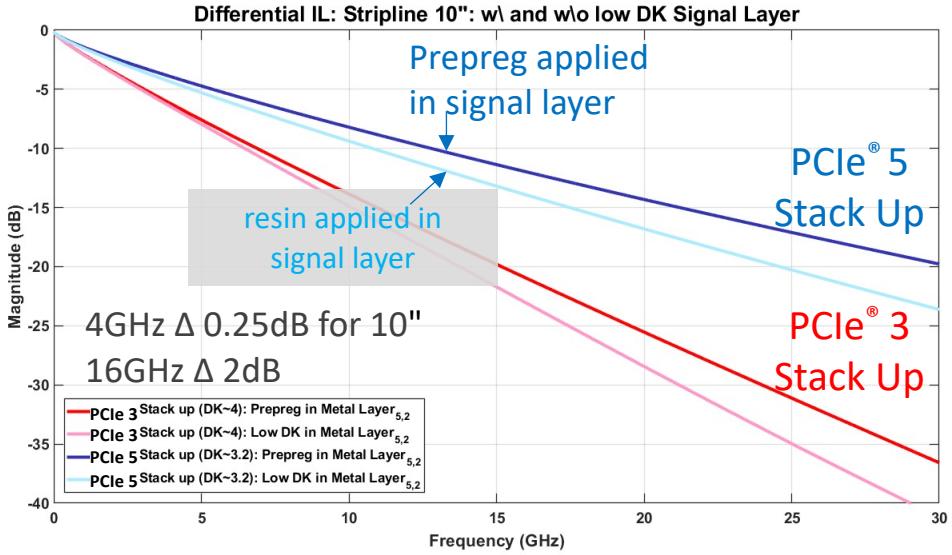

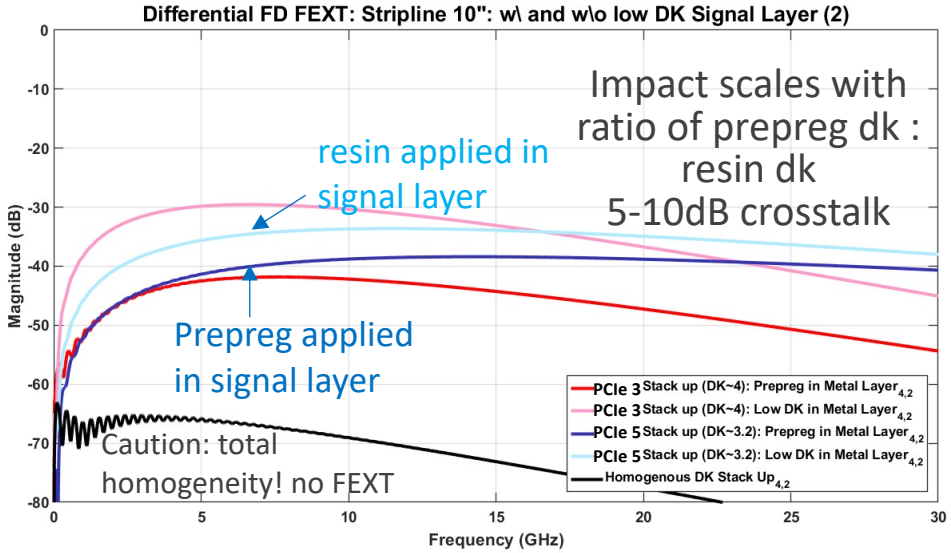

Transmission Lines: Environmental

Temperature and Humidity

- Neglected environment in models or lab may lead to surprises
- Up to 30% in worst case measurement; consider 10-20% in modeling
- Expect localized hot spots
- Microstrip is worst

Nominal Loss (dB/inch at 4 GHz)	Microstrip Worst Case 75°C dB/inch (percentage increase)	Stripline Worst Case 75°C dB/inch (percentage increase)
0.40	0.50 (25%)	0.48 (19%)
0.44	0.54 (23%)	0.52 (17%)
0.5	0.60 (20%)	0.58 (15%)
0.6	0.70 (17%)	0.68 (13%)
0.7	0.93 (32%)	0.90 (29%)
0.8	1.03 (28%)	1.00 (25%)

Note: *Decibels per inch (dB/inch) is measured at 4 GHz.*

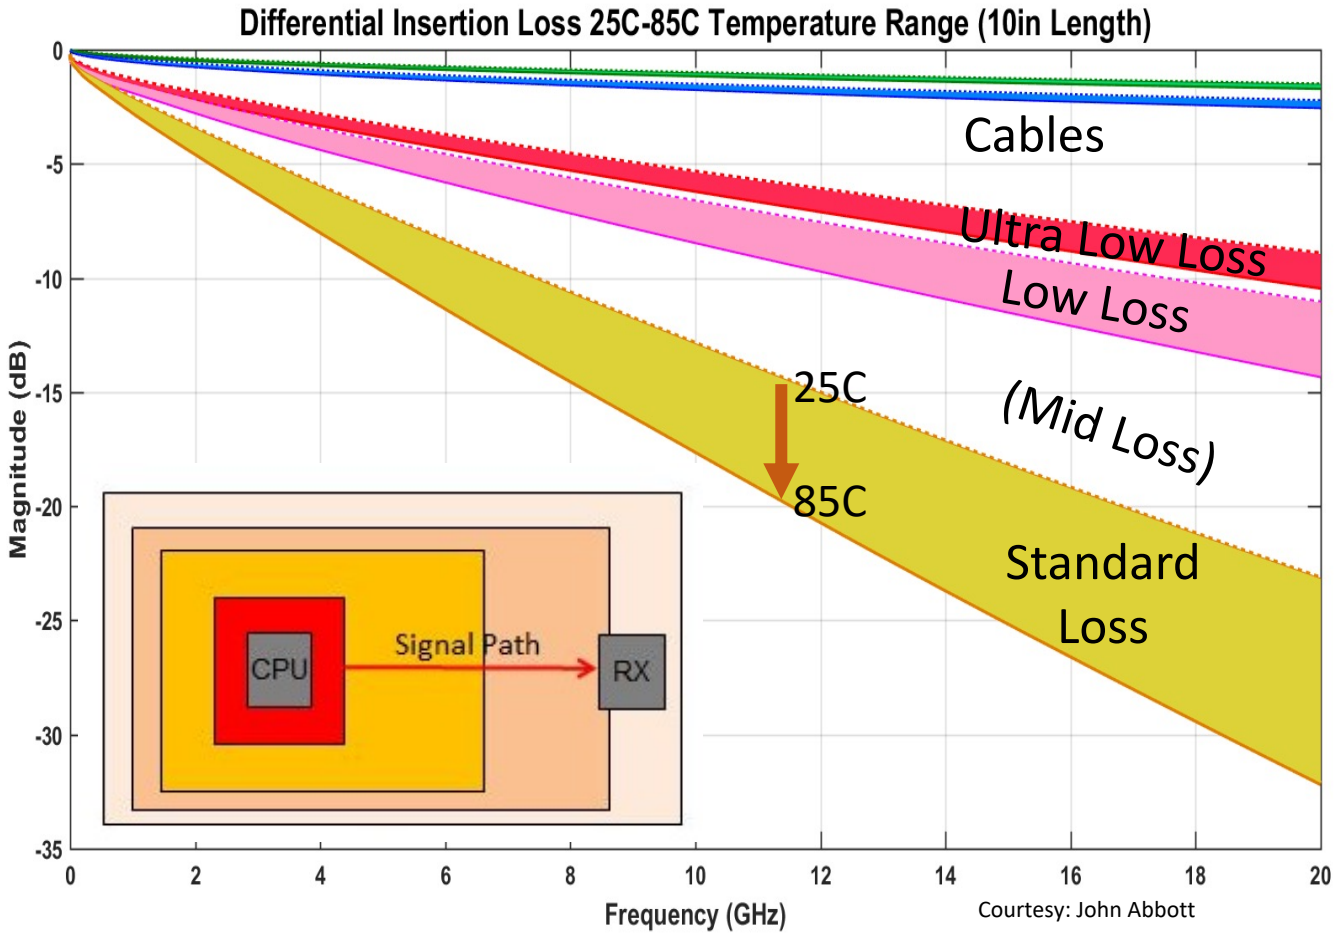

Courtesy: John Abbott

Loyer J, Kunze R, Brist G. Humidity and Temperature Effects on PCB Insertion Loss, DesignCon 2013.
Zhang H, Ou J, Krooswyk S. High Speed Digital Design - Design of High Speed Interconnects and Signaling

Via Design Considerations: Impedance

- First order optimizations
  - Closely coupled pair and oval anti-pads
- Second order
  - Uniquely tuned anti-pad by-layer
- Lastly,
  - Laser\CCD drill, back drill
  - Drill and pad diameter reductions

First Order Optimizations:

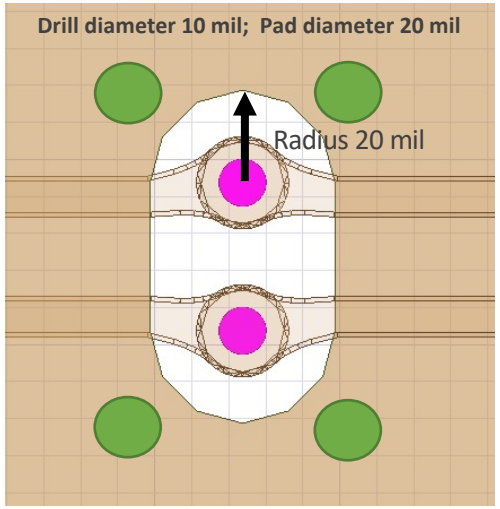

Second Order:

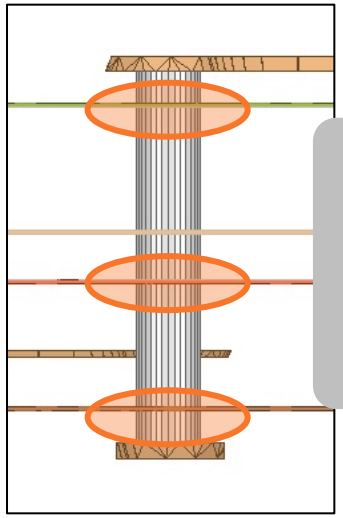

Unique and larger antipads on planes adjacent to pads

And\or backdrills

Don't forget the grounds

Return Path Crosstalk

Should we use simulation to optimize the via? If we do not, how much stub is tolerable? (10/20/30 drill/pad/anti)	
PCIe® 3	Not necessary, stubs <=80mils no problems
PCIe® 4	Optimization may help via stubs 50 mils and larger
PCIe® 5	Expect optimization simulation for all via designs

Via Design Considerations: Stub vs. Xtalk

- Two via designs considered – what's best?
  - We often neglect the crosstalk in the shorter via stub
  - Consider SNR to make decisions, e.g., ICR
- More significant issue in space-constrained BGA fields where back drill may not be possible
  - Greater concern for NEXT

Two Breakout Options

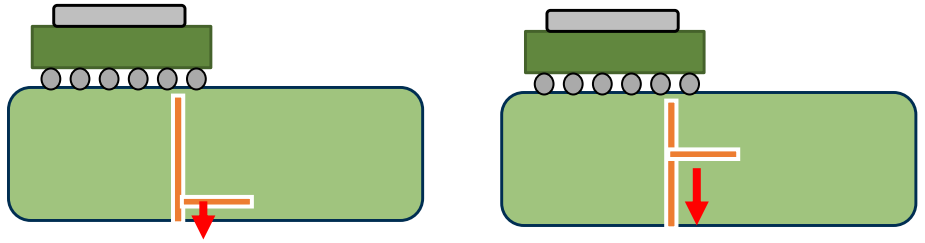

Better Reflection
More Crosstalk

Worse Reflection
Less Crosstalk

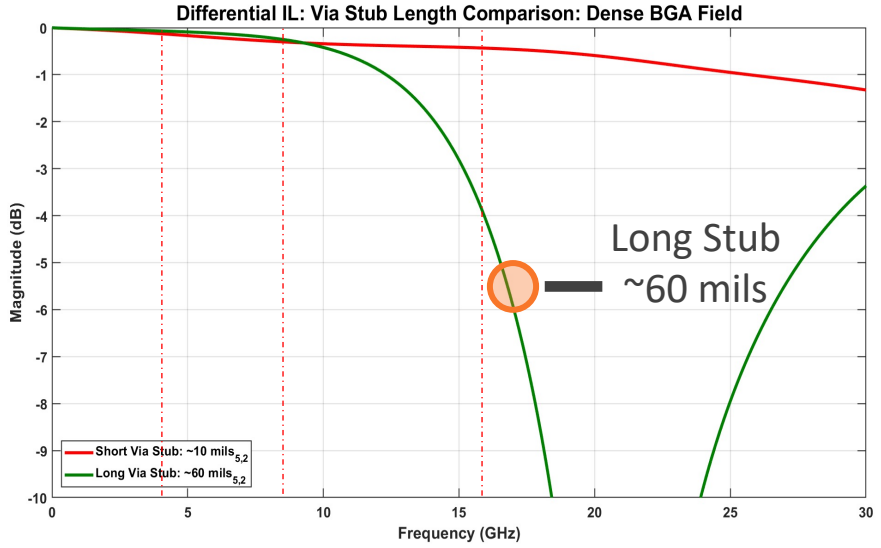

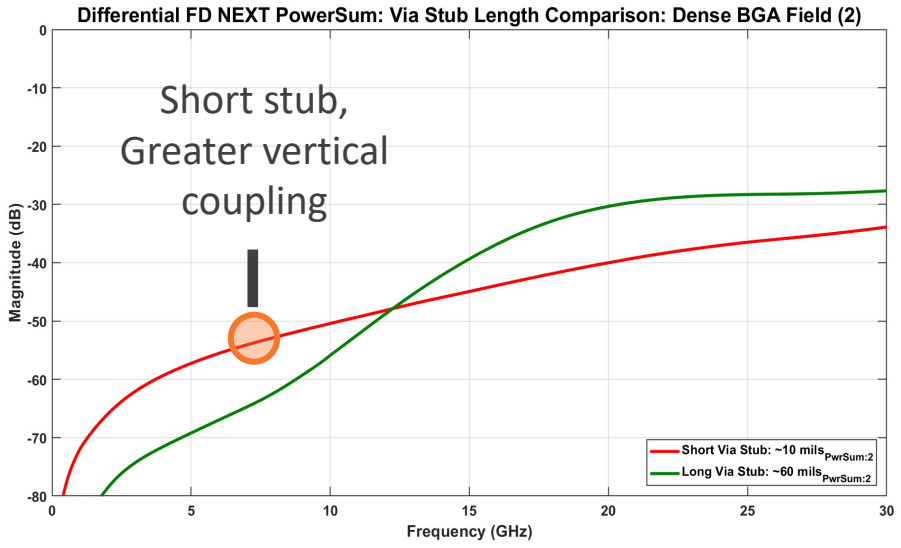

Via Design Considerations: Placement

Boomerang Via

- Addition of small via is negligible compared that large THM\PFT stub on Gen3-4 designs
- No additional via for capacitor!

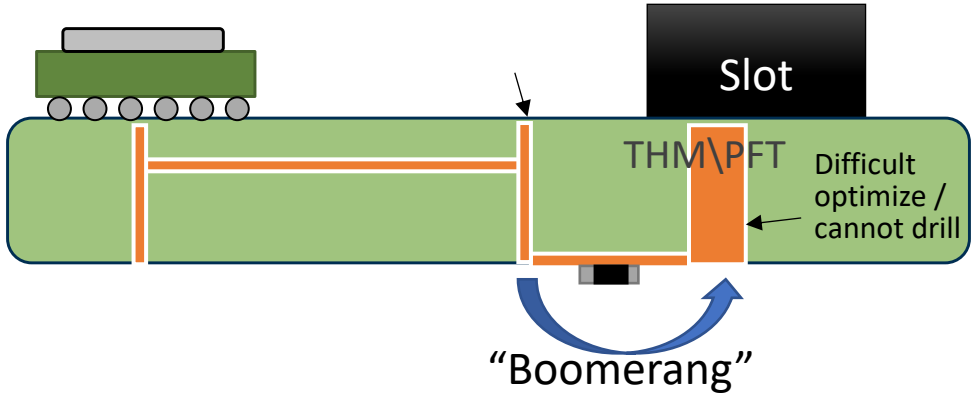

Keep Surface Transitions Short

- Keep lengths to 0.5 inch or less w\o caps
- Crosstalk + independent Z variable

Surface Breakout Advantages:

- Mid-channel vias are more easily optimized
- Easy capacitor placement

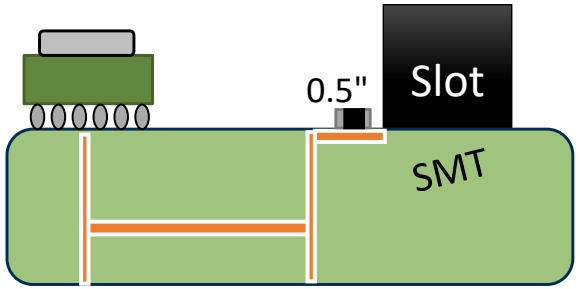

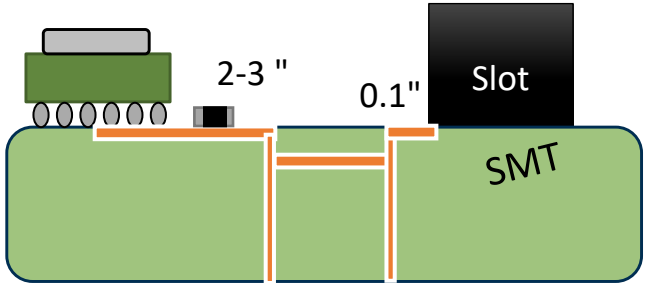

Connector Selection: CEM Interface

- Move to Surface Mount at Higher Speeds
  - At Gen4, some TH\PFT may still be possible
- FD Limits for CEM Interface
  - IL, RL, Crosstalk
- New Crosstalk Compliance Approaches at PCIe® 5.0
  - Recognize pin maps in FEXT+ NEXT Power Sums
  - Compliance by integrated crosstalk

Version	Available Connectors
PCIe® 3.0	SMT or TH\PFT
PCIe® 4.0	SMT or TH\PFT
PCIe® 5.0	SMT Only

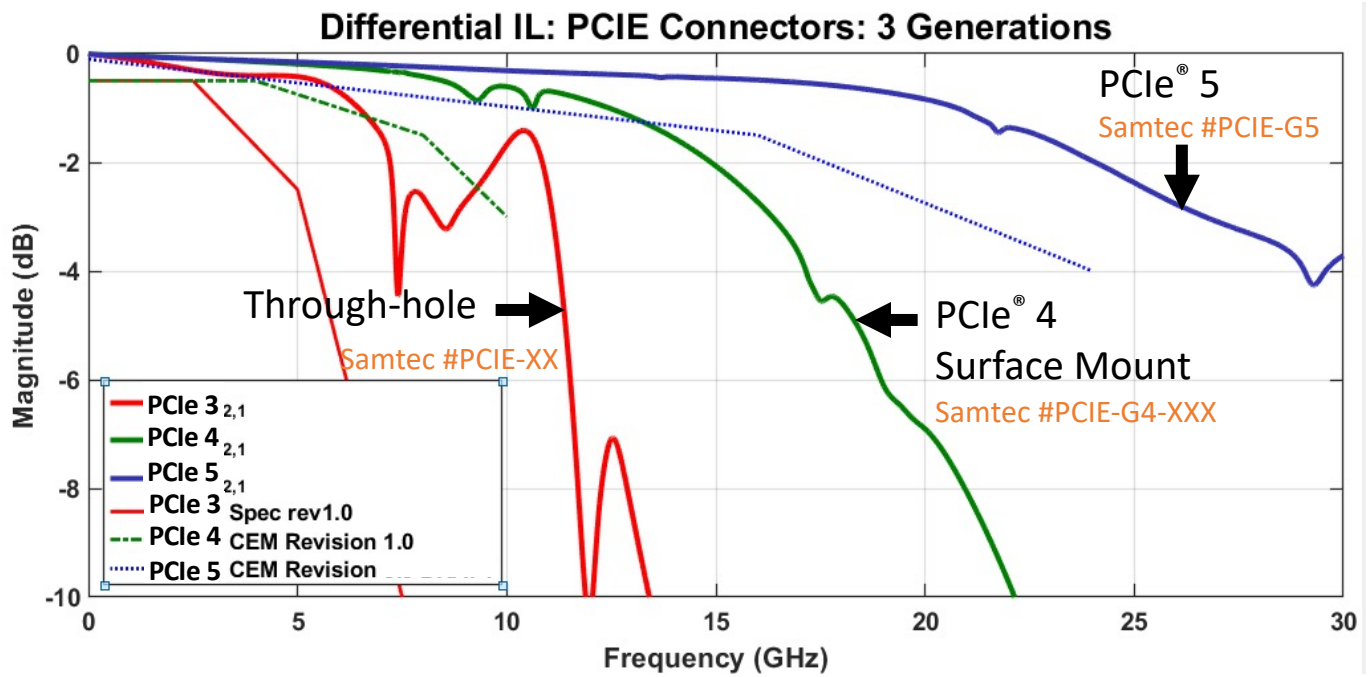

For more details: Latest PCI Express® Card Electro-Mechanical Specification and Gen5 Connector Compliance with Integrated Crosstalk Noise, S. Krooswyk, PCI Express® Developers Conference October 2020

Connector Selection: Non-CEM \ Captive

- Specification guidance is validation thru channel compliance (simulated eye height and width)

  😂 This feels a lot of burden, why don't CEM FD limits apply?

- Other form factors are taller – having more loss and lower resonances, with more pins and noise

- Anecdotal recommendations:

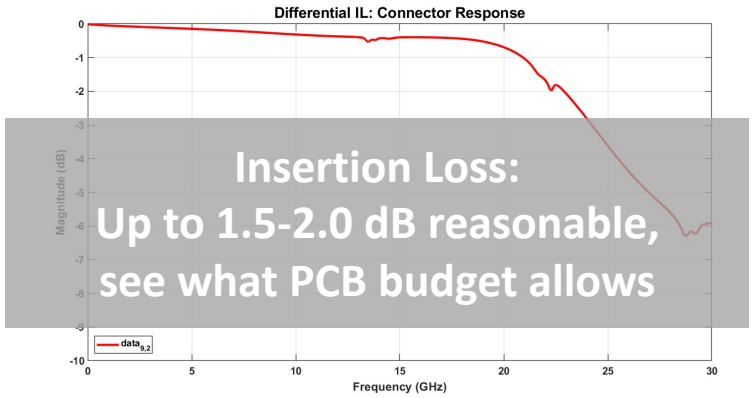

Insertion Loss:
Up to 1.5-2.0 dB reasonable, see what PCB budget allows

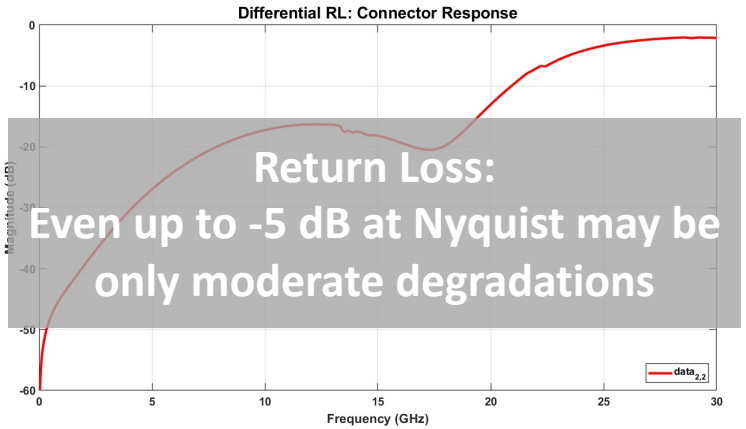

Return Loss:
Even up to -5 dB at Nyquist may be only moderate degradations

Examples: Don't Judge a Bit Just by its Fourier – 112 Gbps PAM4 Component Optimization and Selection, DesignCon 2019, S. Krooswyk, M. Rengarajan

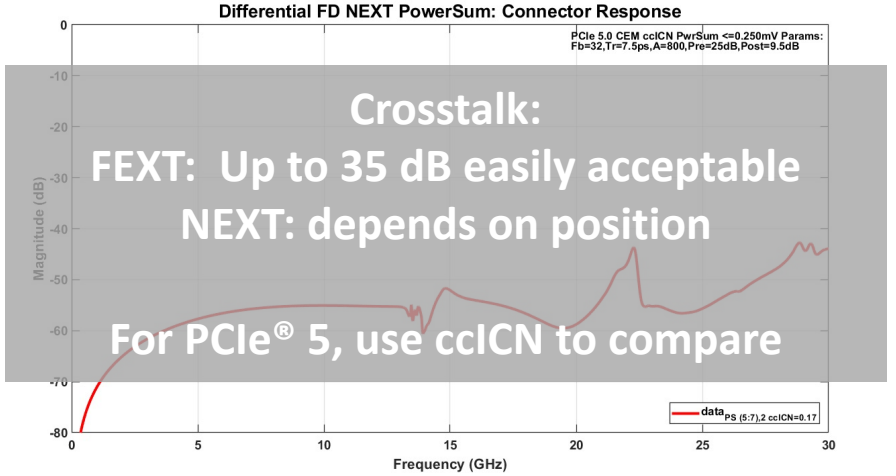

Crosstalk:
FEXT: Up to 35 dB easily acceptable
NEXT: depends on position

For PCIe® 5, use ccICN to compare

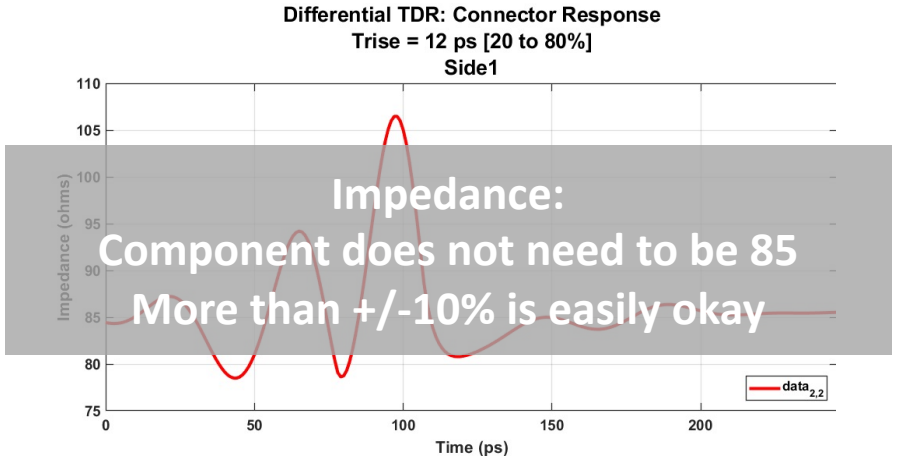

Impedance:
Component does not need to be 85
More than +/-10% is easily okay

gEEk® spEEk April 30, 2020 – PCI Express®: is 85 Ohms Really Needed?

Cable Considerations

- Cable Loss Advantage for Longer Reach
  - Relatively temp & humidity insensitive
  - High value for all generations

Table: Raw 34 Gauge Cable Conversion at Nyquist	
PCIe® 3.0 Standard Loss PCB	10" PCB : 85" Cable 34 Gauge
PCIe® 4.0 Mid-Loss PCB	10" PCB : 94" Cable 34 Gauge*
PCIe® 5.0 Low-Loss PCB	10" PCB : 55" Cable 34 Gauge

*Smaller increment from Standard to Mid Loss PCB, than Mid to Low 1-inch of 34-Gauge Twinax is 0.20 dB at 16GHz.

- Greater Interest for PCIe® 4.0/5.0 where PCB Solution Space is Reduced
- Avoid PCBs for Better Return Loss

- Positions for side bands – know your features 2-8 pins per port
- Optimal impedance?
  - the mating connector is the greatest factor
  - lower losses in 100 ohm twinax
  - EH\EW simulations better @ 100 ohm
  - see… gEEk® spEEk April 30, 2020 – PCI Express®: is 85 Ohms Really Needed?

Cable Reach Example: PCIe® 5.0

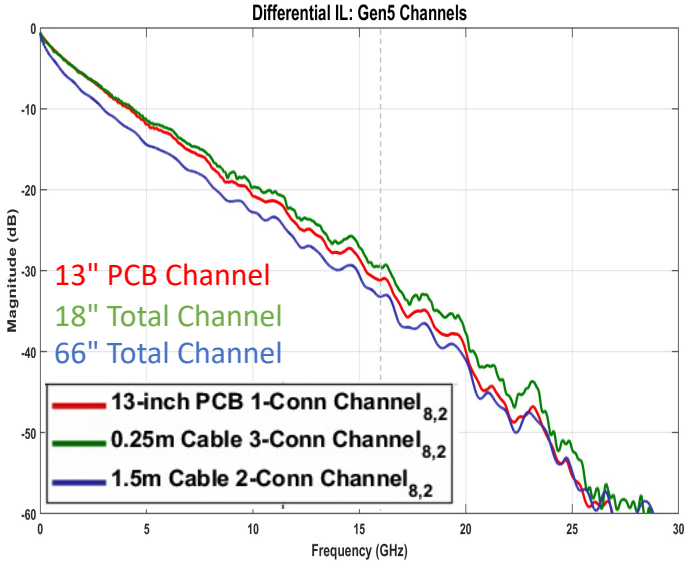

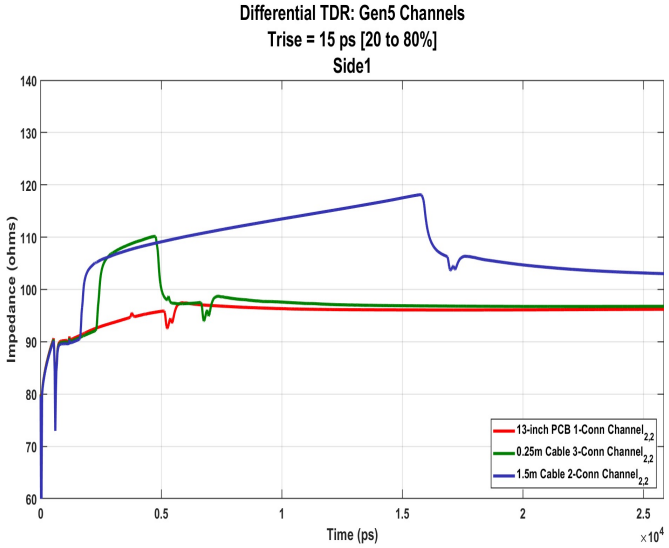

13" PCB 1-conn:
31.1mV/0.48UI

0.25m Cable, 3-conn:
32.9mV/0.5UI

1.5m Cable, 2-conn:
32.4mV/0.6UI

Spec Min: 15mV/0.3UI

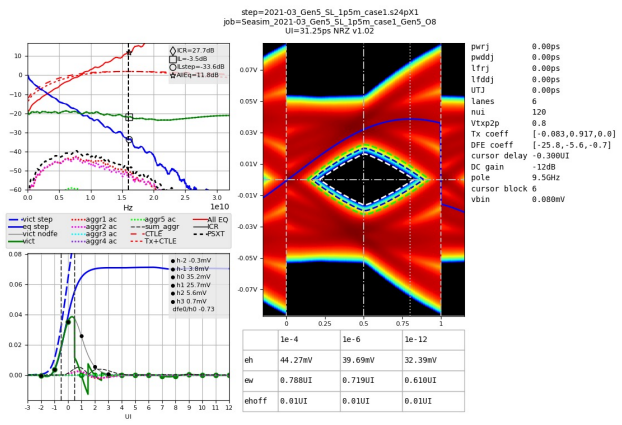

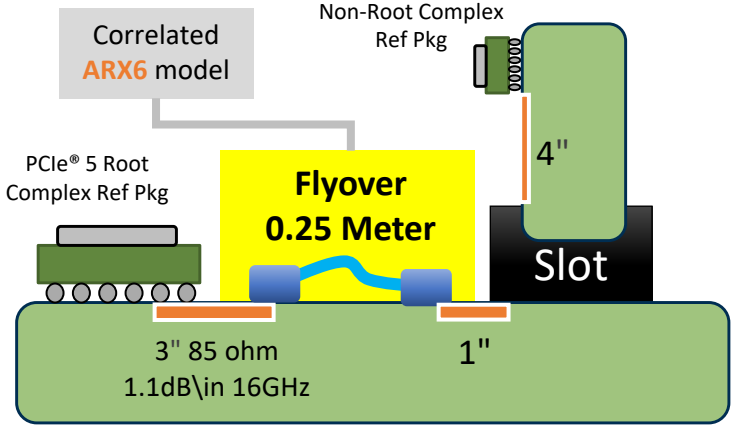

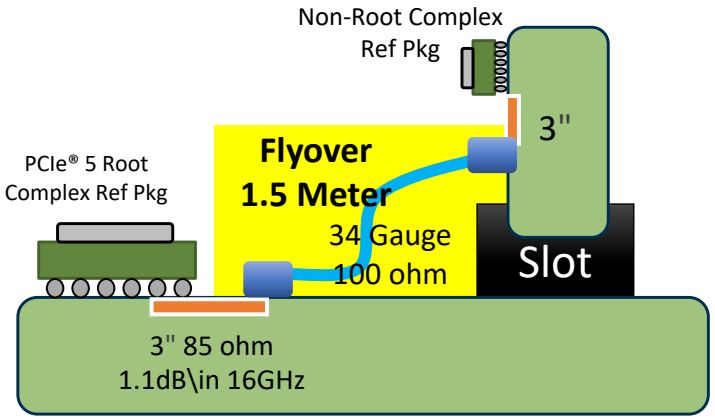

Stripline 3x separation crosstalk models; PCIe® 5 via not included.

Channel Analysis Steps: Top-Down Approach

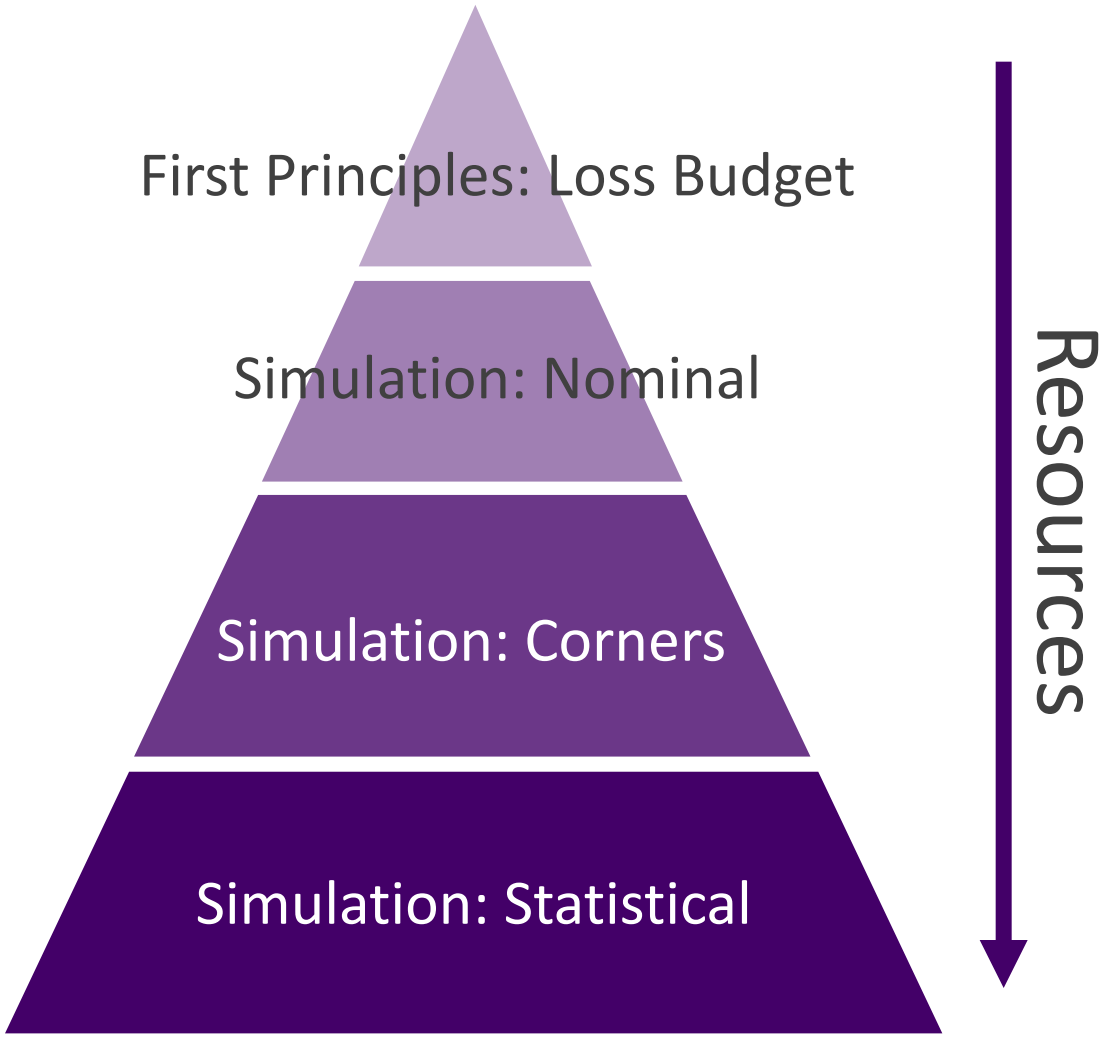

- At the top – easiest analysis first
- Step down the pyramid to mitigate risk
- Increasing effort towards bottom

First Principles: Loss Budget

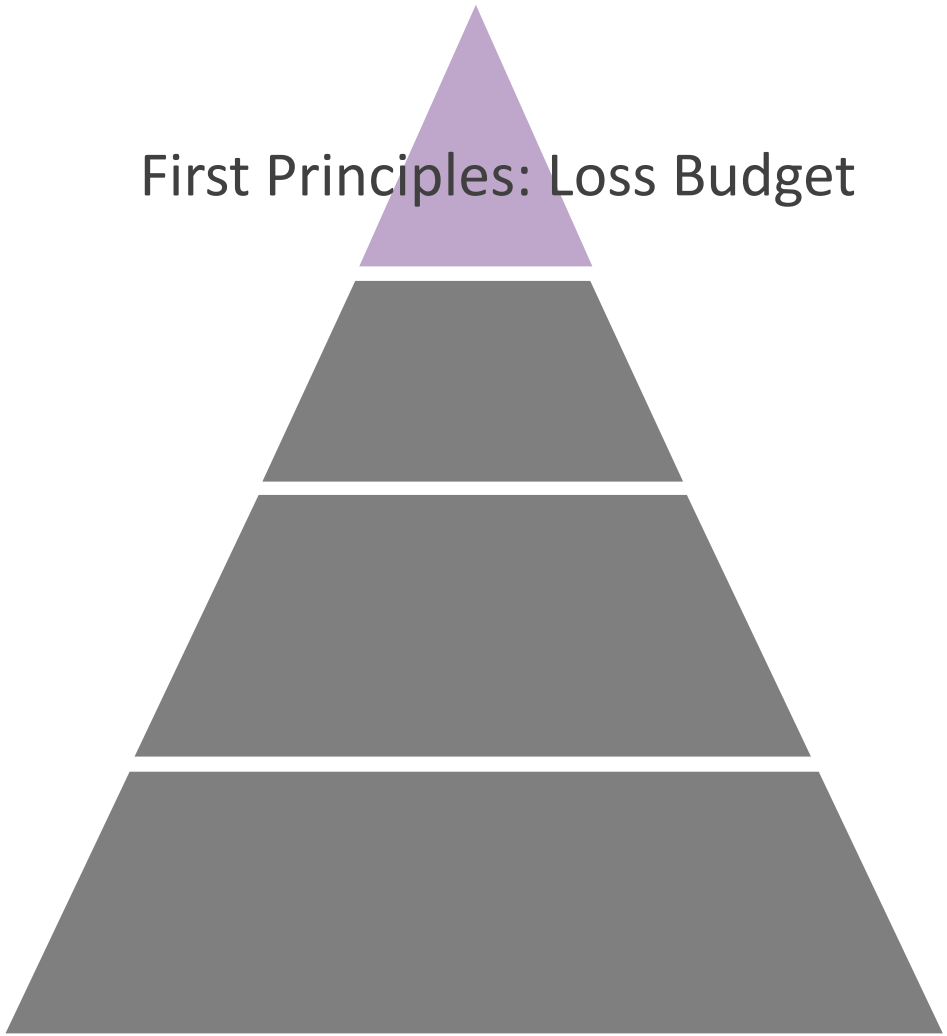

- Initial pass at channel compliance
- No simulation needed
- Spreadsheet based calculations
- Stop gap to identify channels with too much insertion loss
- Specification is not normative about loss – common to see PCIe® 3.0/4.0 channel simulation complaint slightly *beyond* loss target

First Principles: Loss Budget Example

- Reasonable budget calculations for all three generations
  - Validated with simulation
- Package losses per spec
- Card losses are informative
- At PCIe® 4.0 and above, a remainder must exist for reflection and crosstalk noise missed on paper

	PCIe® 3.0	PCIe® 4.0	PCIe® 5.0
Target	23.5 dB	28 dB	36 dB
PCB Main Board	16"	10"	10"
	0.8 db\inch	1.35 db\inch	1.1 db\inch
	@ 4 GHz	@ 8 GHz	@ 16 GHz
	12.8	13.5	11
RC Package	3.5	5	8.5
Vias, Caps	0.5	1	1.5
Connector (Spec)	0.5	0.5	1.5
4" Card Budget	6.5	8.5	9.5
Reflection & Crosstalk Penalty	negligible	< 1	4
Total	**23.8 dB**	**28.5 dB**	**36 dB**

Channel Compliance Simulations

- Defined by PCI Express® Base Specification
  - Any simulation tool that adheres to requirements is acceptable
  - Phyton tool is provided to PCI-SIG members
- A compliant channel passes eye diagram in all impedance and environment corner conditions
- Reference packages are minimally compliant to loss and return loss
  - Package model crosstalk is a reference, and not a requirement
- Refine solution with device vendor models
- Channel compliance is not actually measurable, except for CEM slots
  - Captive channels may use vendor determined validation limits for on-die eye or BER

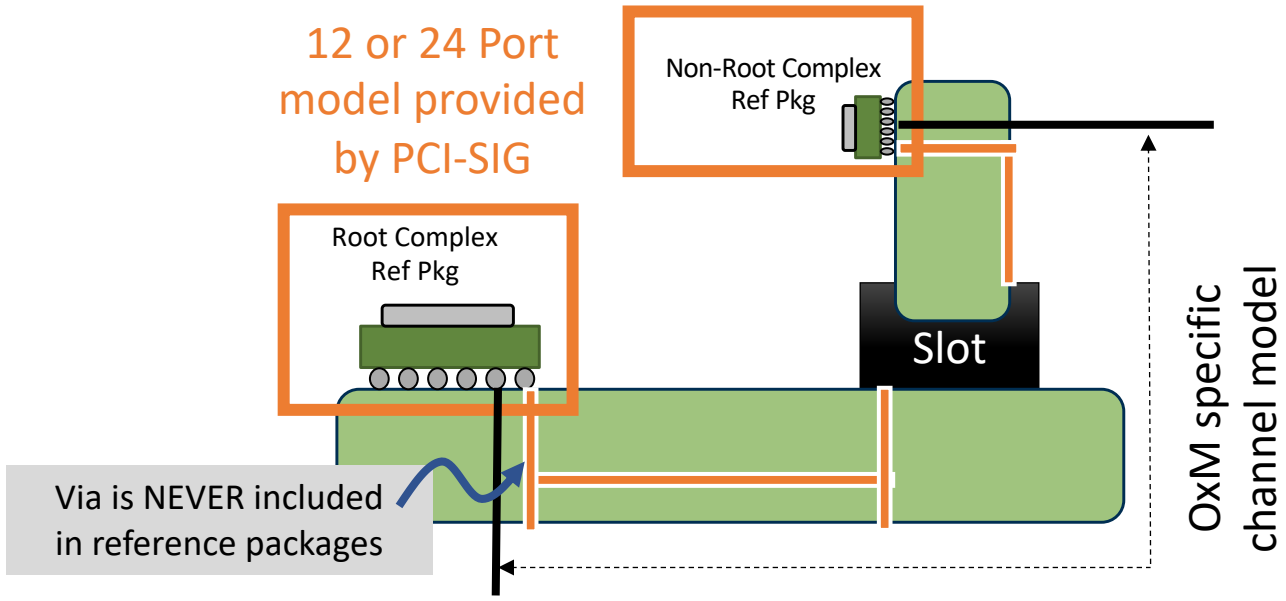

Interpreting Channel Simulation Results

This plot, FD conversion created from time domain, may show small discrepancy from s-parameters.

Equalized channel response

Power sum is total of FEXT and NEXT

Loss

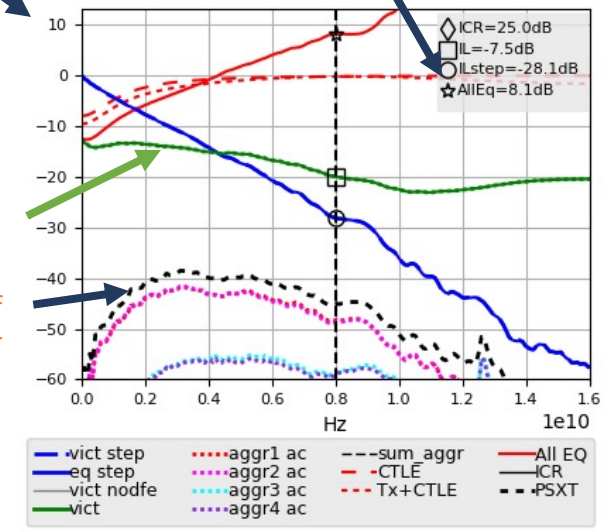

Time domain responses (same colors)

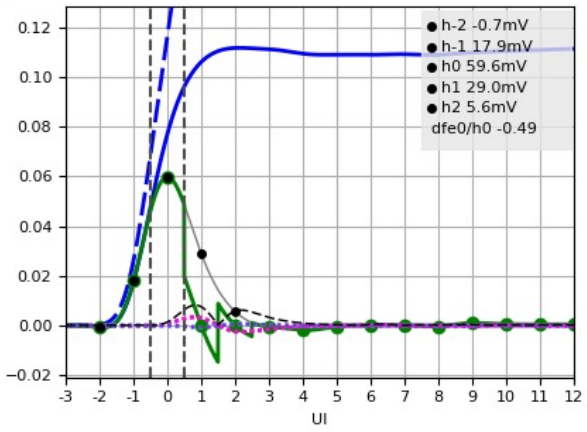

step=2021-03_Gen4_SL_varLen_case4.s20pX1
job=Seasim_2021-03_Gen4_SL_varLen_case4_Gen4_O7
UI=62.50ps NRZ v1.02

Data Rate

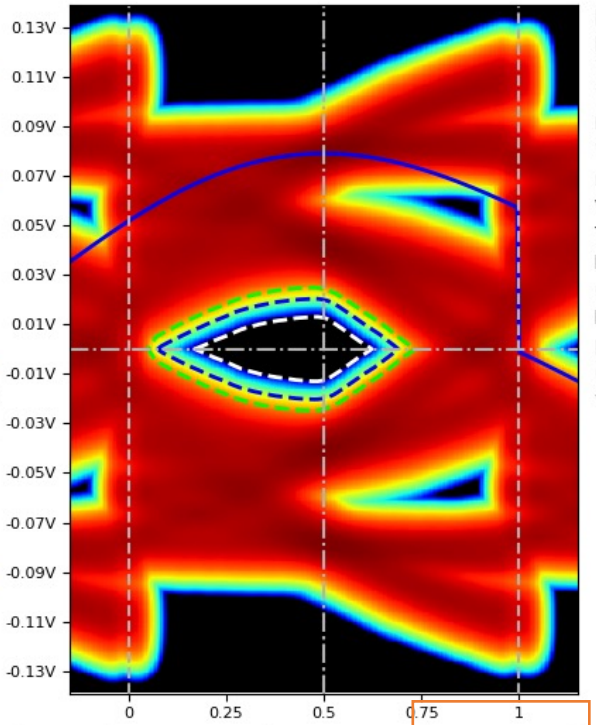

pwrj	0.54ps
pwddj	5.00ps
lfrj	0.74ps
lfddj	3.75ps
UTJ	17.96ps
lanes	5
nui	120
Vtxp2p	0.8
Tx coeff	[-0.042,0.917,-0.042]
DFE coeff	[-29.0,-5.5]
cursor delay	0.000UI
DC gain	-8dB
pole	2.0GHz
cursor block	6
vbin	0.500mV

Total: Victim + Aggressors

Adaptive EQ Results

	1e-4	1e-6	1e-12
eh	49.58mV	40.62mV	26.02mV
ew	0.683UI	0.615UI	0.453UI
ehoff	-0.03UI	-0.03UI	-0.02UI

Total Eye Opening, peak to peak, Compare to spec

EH Offset: location of maximum eye height reported value

Channel Simulation Results: Example Length Sweep, PCIe® 4.0

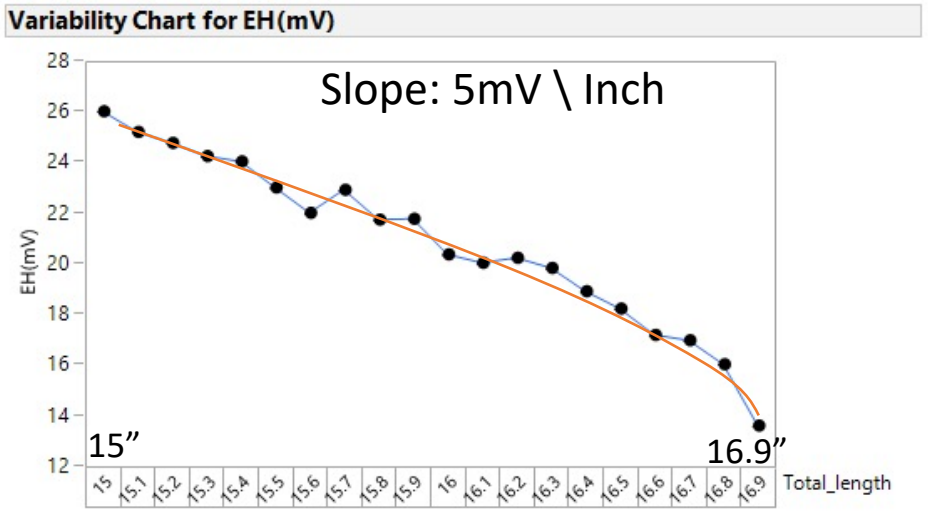

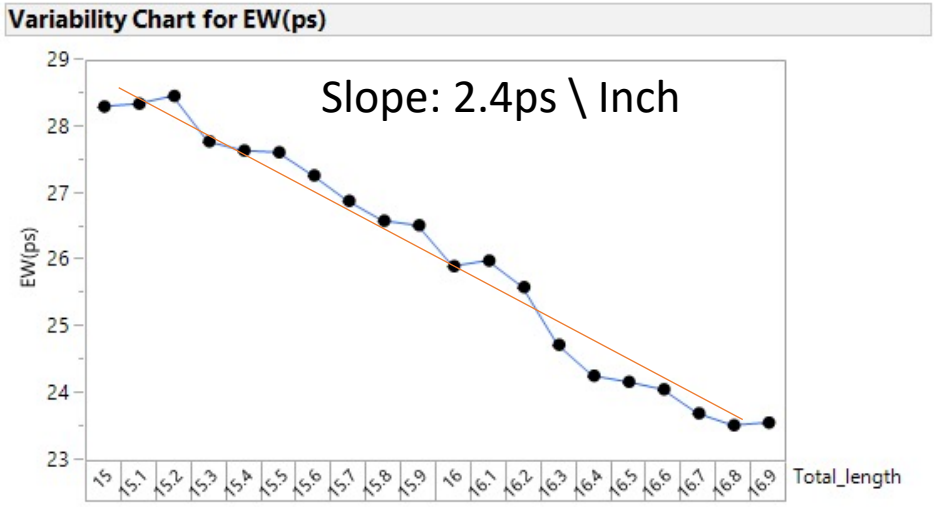

- Some dithering occurs in simulation results
  - Example: near +/- 1mv and 0.5 ps deviation from a smooth curve
- Sources:
  - Discrete steps in adaptive equalization
  - Reflection alignment with CDR
- Caution simulating a single-point for a small effect: could the conclusion be in the noise?
- From simulation, we can determine:

PCIe® 4.0 Rule of Thumb

5 mV \ 2.4 ps = 1 inch = 1.4 dB at 8 GHz

Channel Simulation Results:
Eye Closure Slope \ Rule of Thumb

- What to do with this information?
- Apply to simulation result comparisons when asking “What is the impact”
- Translate confusing “mV impact” into meaningful loss or PCB terms
- Clarity will be given to negligible effects

PCIe® 3.0 Rule of Thumb
4mv \ 1ps = 1 inch = 0.8 dB at 4 GHz

PCIe® 4.0 Rule of Thumb
5mV \ 2.4 ps = 1 inch = 1.4 dB at 8 GHz

PCIe® 5.0 Rule of Thumb
5mv \ 0.5 ps = 1 inch = 1.1 dB at 16 GHz

Simplify Decisions:
e.g. “The cost is X mV / X ps, which is a X-inch reduction”

Sensitivity to Neck Down Regions

- Tight regions are inevitable
- If regions are FEXT (and most are), there is no location sensitivity
  - Occurrence near Tx, or near Rx, is no different
- Short regions of tight spacing are acceptable for shorter durations
- Let's review two excursions:
  - Stripline at 2x dielectric thickness
  - Microstrip at 5x dielectric thickness

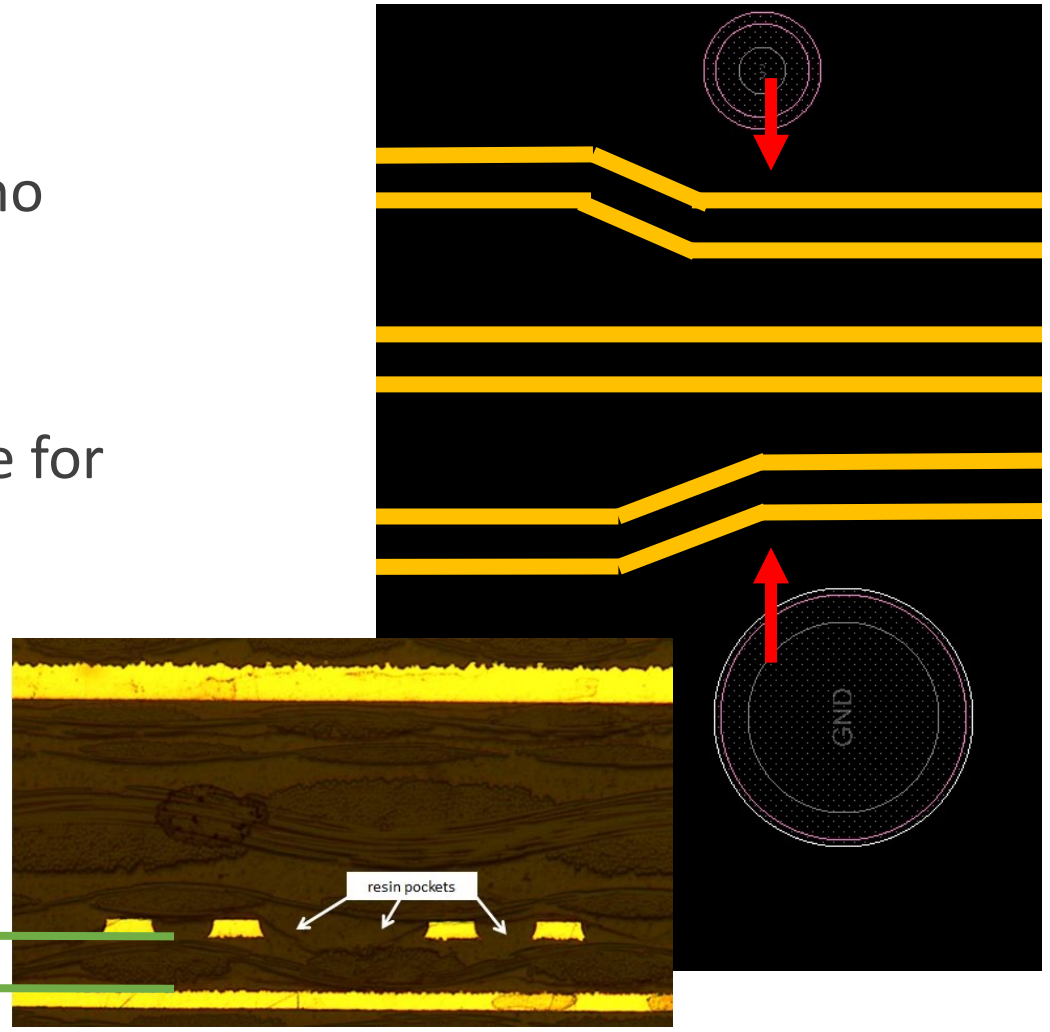

Dielectric thickness is taken for nearest reference plane.

Sensitivity to Neck Down Regions: Stripline, by PCIe® Version

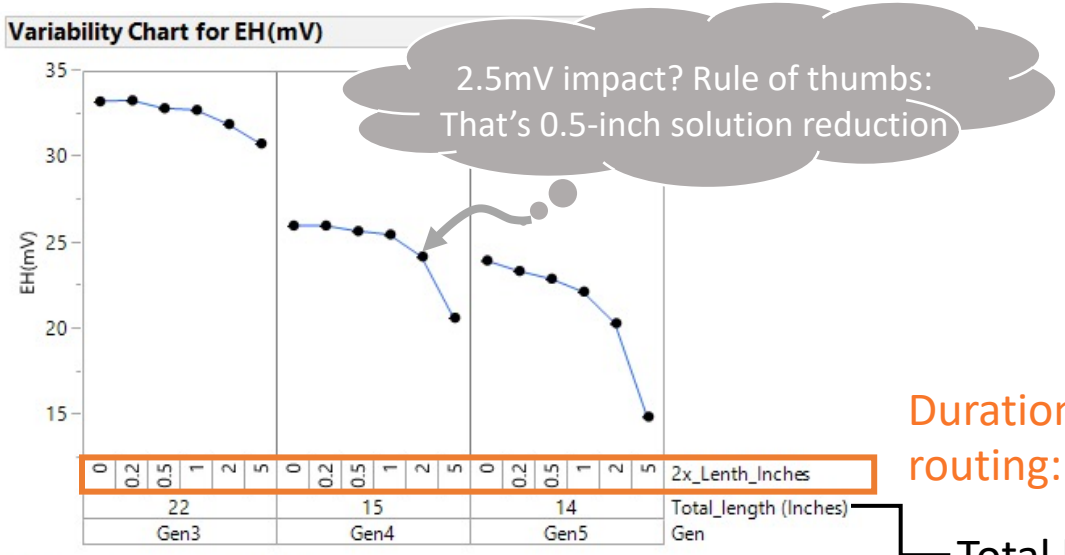

2.5mV impact? Rule of thumbs: That's 0.5-inch solution reduction

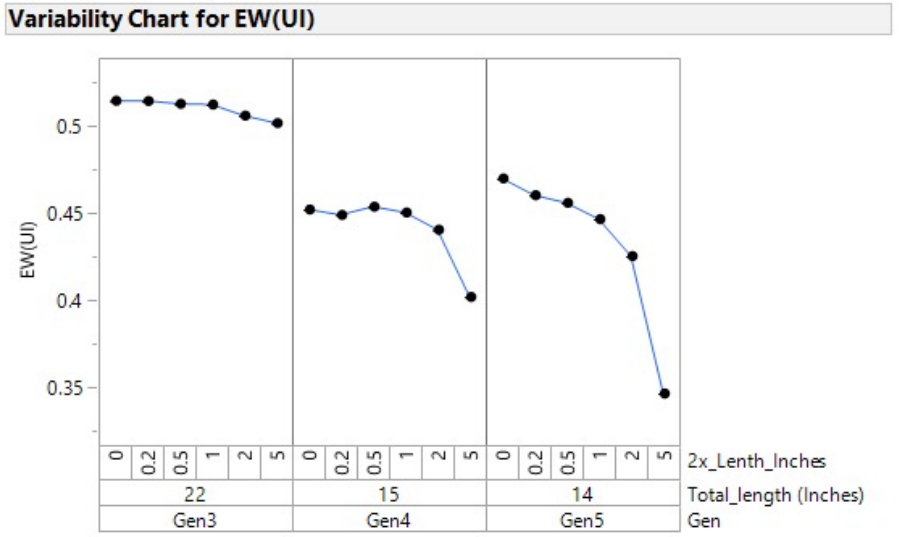

Duration of 2x neck down routing: 0, 0.2, 0.5, 1, 2, 5"

Total length constant

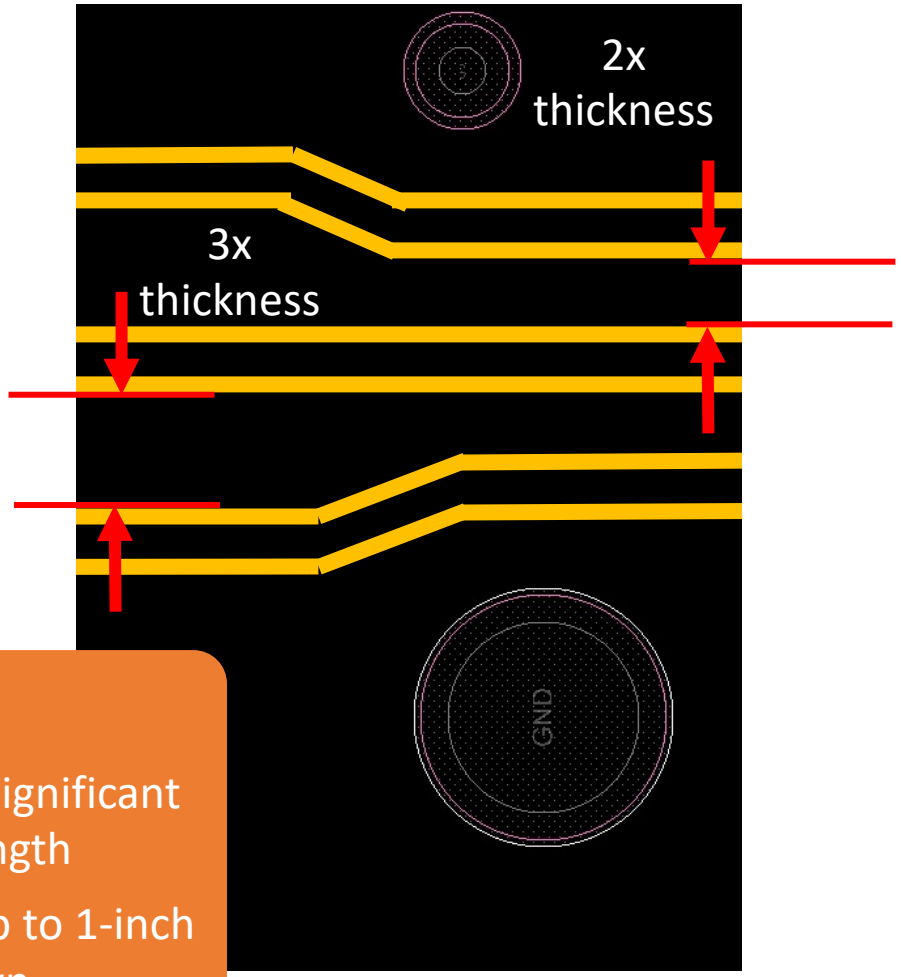

TAKE AWAY

PCIe® 3 is agnostic to significant 2x neck down length

PCIe® 4 is insensitive up to 1-inch of 2x neck down

Sensitivity to Neck Down Regions: Microstrip, by Generation

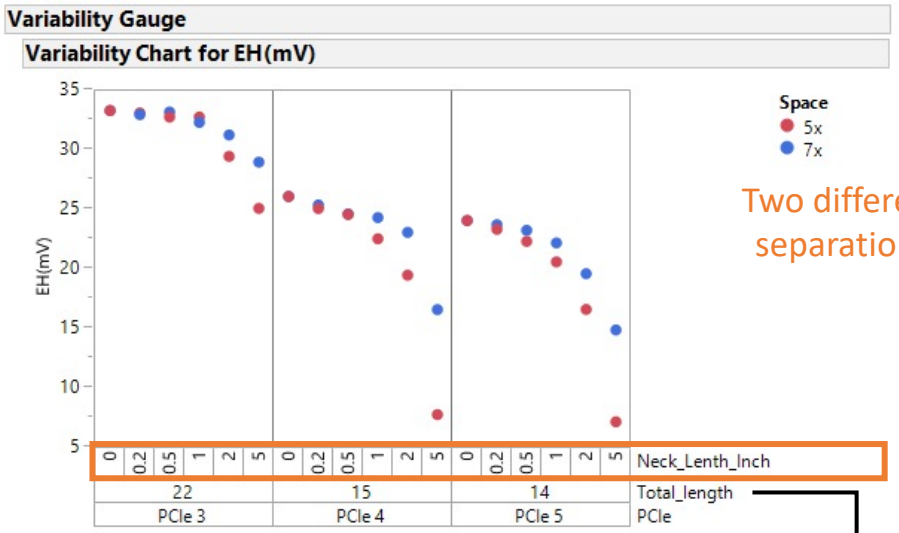

Two different separations

Duration of 5-7x neck down routing: 0, 0.2, 0.5, 1, 2, 5"

Total length constant

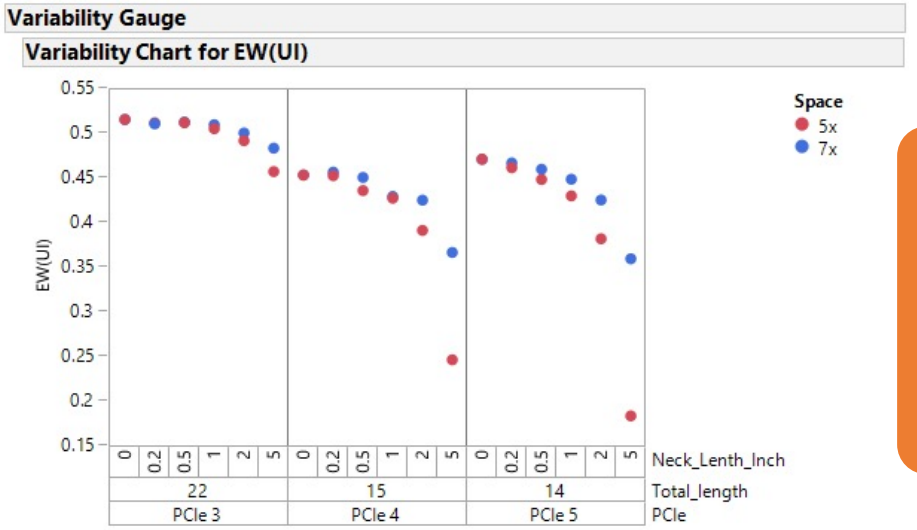

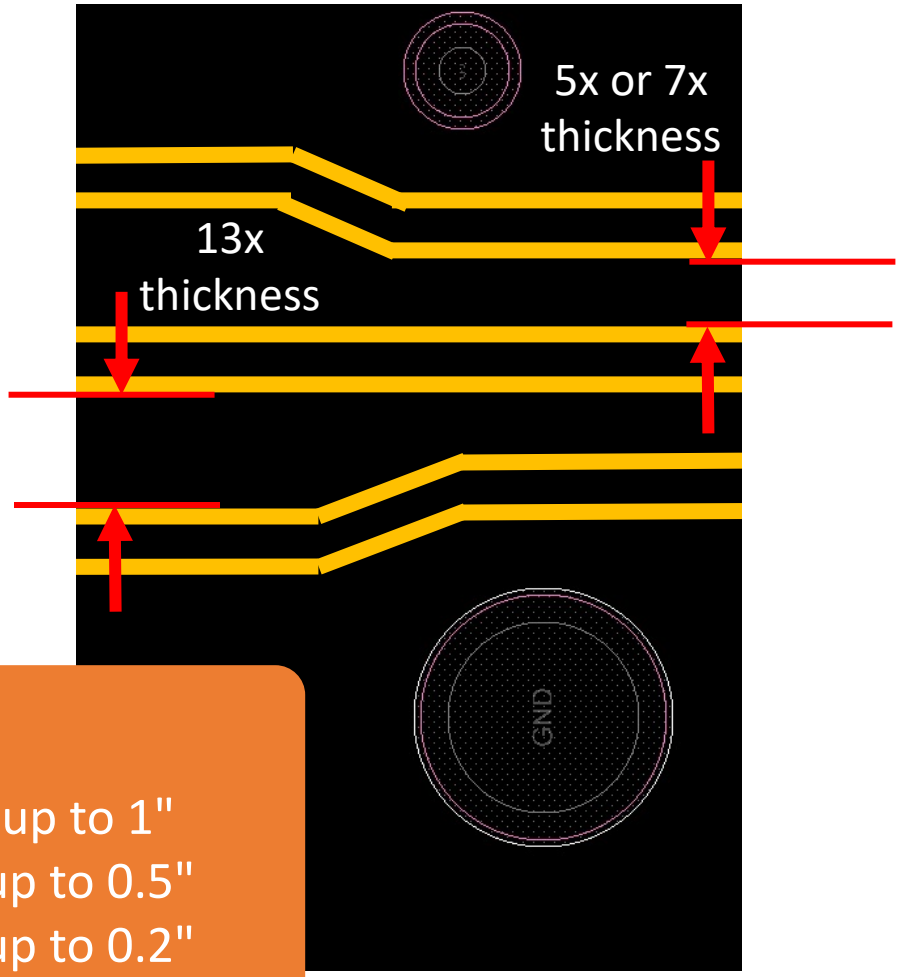

TAKE AWAY

PCIe® 3: 5x negligible up to 1"
PCIe® 4: 5x negligible up to 0.5"
PCIe® 5: 5x negligible up to 0.2"

Overall Topology Guidelines

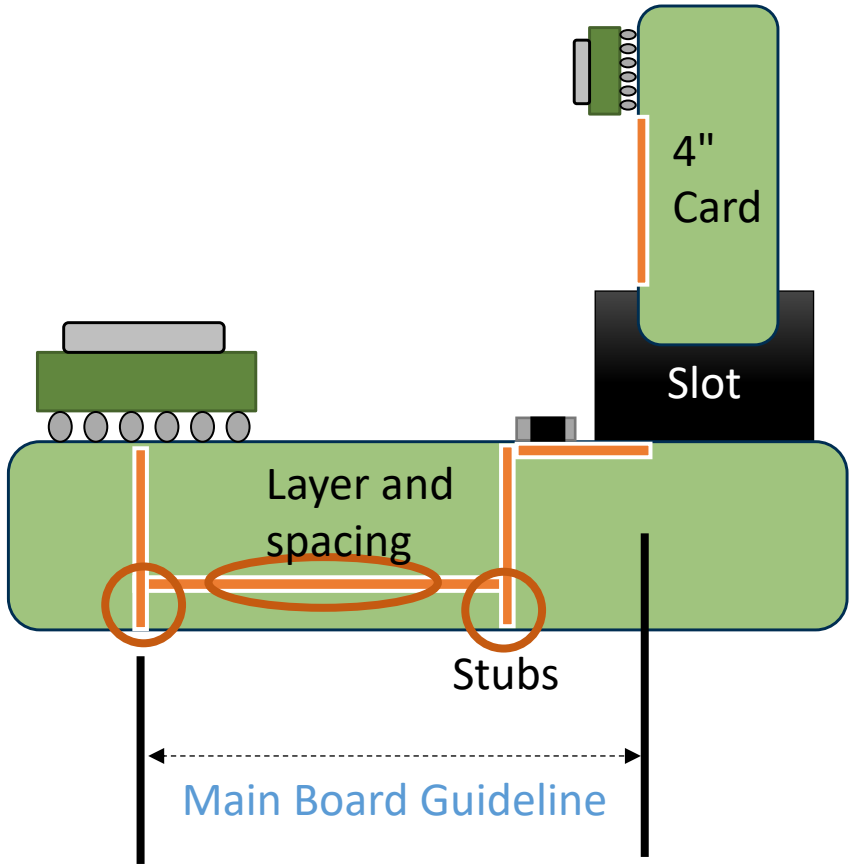

	PCIe® 3.0	PCIe® 4.0	PCIe® 5.0
Option 1 (Stripline)	16" .8db\in Stub: <= 80 mil PFT Conn Stub: <= 80 3x Stripline Space	10-12" 1.4db\in Stub: <= 60 mil SMT Conn 3x Stripline Space	10" 1.0db\in Stub: managed 3x Stripline Space
Option 2 (Mixed)	16" .85db\in Stubs: <= 80 mil 9x Microstrip Space	9-10" 1.4db\in Stubs: <= 60 mil 11x Microstrip 7-8" 9x Microstrip	8-9" 1.0db\in Stubs: managed 11x Microstrip

Repeater Considerations: Two Types

Re-Timers

- A2D, D2A Conversion
- Re-Clocked
- Up to 2x channel reach
- Full channel capability available before AND after
- Specification supporting 2 re-timers for PCIe® 4.0
- Reference pinmap

Re-Drivers

- Analog equalizer configs validated in lab
- ~30% channel reach increase
- Ideal for captive channels
- May not support open interchangeable slots
- Adaptive TXEQ phases disabled

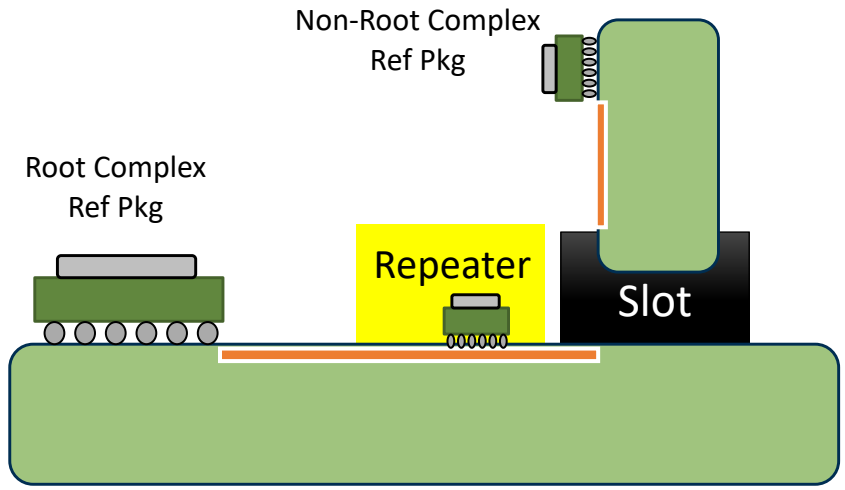

PCI Express® Retimers vs. Redrivers: An Eye-Popping Difference, Casey Morrison https://pcisig.com/pci-express%C2%AE-retimers-vs-redrivers-eye-popping-difference
High Speed Series Bus Repeater Primer, Samie Samaan et al https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/serial-bus-white-paper.pdf

Summary

- As expected, greater efforts are needed with higher speed, however the challenges are not insurmountable
  - With hope, this tutorial better equips designers
- Some reminders…
- Critical considerations for transmission lines include roughness and the environment
- PCB, cable, and connector impedance need not always be 85 ohms
- PCB: cable physical displacement is significant for all generations
- Repeaters include re-timers and re-drivers, for which there are many differences
- Non-CEM connectors have no guideline – so we've offered one
- Channel simulation EH\EW is not always intuitive – decisions are best communicated as PCB-reach or loss

gEEk® spEEk


samtec.com/geekspeek


geekspeek@samtec.com


SIG@samtec.com