PCI Express® 3.0 PHY Electrical Layer Requirements

Dan Froelich

Intel Corporation

Agenda

- PHY Requirements
- Preliminary Jitter Budget
- Statistical Simulation Tools
- 3.0 PHY Rate
- Transmitter Specification
  - ✓ PLL Bandwidth
  - ✓ Reference Location
  - ✓ Timing Parameters
  - ✓ Equalization
- Reference Clock Specification
- Receiver Specification
- Major Form Factor Work Areas
- Next Steps

PCIe® 3.0 Electrical Requirements

- Backwards Compatibility
  - ✓ Gen1/Gen2 cards must operate in Gen3 slots at Gen1/Gen2 performance
  - ✓ 2.0 clocking architectures must be supported.
- Compatible with 2.0 Power Budgets
  - ✓ Low PHY Power Consumption
- Cost: No required changes to connectors, clocks, materials, HVM manufacturing practices.
  - ✓ Extreme server channels may require channel optimizations.
- BER of E-12 or better.
- At least 2x effective data rate of PCIe 2.0 (5.0 GT/s)
- Channel Length Support
  - ✓ Client
    - 1 Connecter, 14” end to end, microstrip, FR4.
  - ✓ Server
    - 2 Connector, 20” end to end, stripline, FR4.

System Jitter Budget 8.0 GT/s

Jitter Contribution	Max Dj (ps) 5.0 GT/s	Max RJ (ps RMS) 5.0 GT/s	Max Dj (ps) 8.0 GT/s	Max RJ (ps RMS) 8.0 GT/s
TX	30	1.4	7	1.6
Ref Clock	0	3.1	0	1.0
Channel	58	0	N/A*	N/A*
RX	60	1.4	11.8	3.6

***Simluation with Statistical Tool Required To Capture Channel Interactions**

Similar Percentages Assumed at 10 GT/s For Rate Investigation

Rate Selection Process

- Select worst case channels.
  - ✓ Several companies provided channel models for HVM 2.0 client and server systems at length target limits.
- Use statistical simulation tools
- Analyze rates that can provide ~ 2x data throughput increase
  - ✓ 8 GT/s with scrambling.
  - ✓ 10 GT/s with 8b/10b.
- Analyze different receiver equalization methods
  - ✓ CTLE
  - ✓ DFE

Statistical Simulation Tools

- Provides jitter relief by moving jitter from Dj bin to Rj bin
  - ✓ For a given channel, enables I/O designers to determine what type, order and equalization resolution is required for a BER target
  - ✓ Accurately models high frequency Tx jitter
- Uses statistically weighted data patterns
  - ✓ More accurate, less conservative than PDA
- Operates on pulse response of channel
  - ✓ Comprehends x-talk, ISI, reflections, etc.
- Accurately models both Common Refclk and Data Driven architectures
  - ✓ Accurately models the interaction of CDRs and ISI
  - ✓ Simulates clock models with supply noise sensitivity, device thermal noise, duty-cycle error and jitter amplification

E.g.: Statistical Treatment of Jitter

- Consider T_{MIN_PULSE} parameter
  - ✓ Defined to limit channel induced jitter amplification
- 5.0G spec defines T_{MIN_PULSE} as 0.1 UI (max)
  - ✓ 5.0G spec makes no assumptions regarding Dj/Rj breakdown
  - ✓ This method of budgeting T_{MIN_PULSE} assumes jitter is 100% bimodal Dj
  - ✓ Equivalent to 20 ps Dj, 0 ps Rj
- Analysis of Tx jitter sources yields different results
  - ✓ Jitter over 1.5G – Nyquist will generate jitter amplification
  - ✓ Rj and Dj over this range tend to be spectrally flat
  - ✓ Substantial reduction of Dj can be achieved

Statistical Signaling Analysis

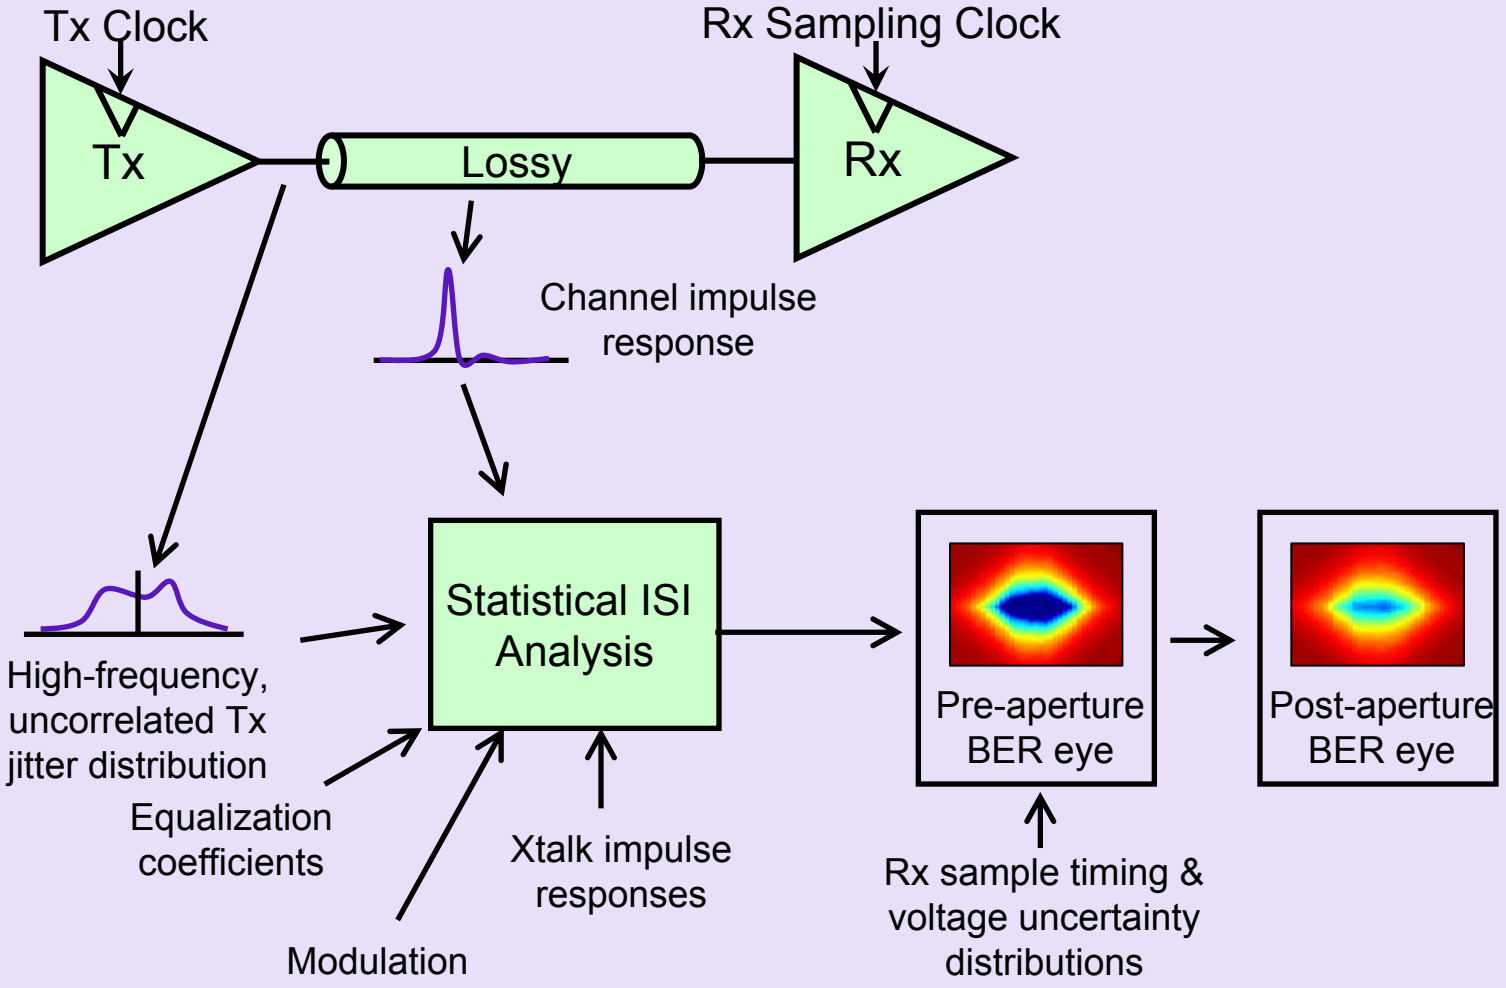

Client Channel Configuration

Seg	Description
A	MCH PKG
B	Break Out
C	MB Main 7”
D	MB post cap
F	Add in card main 3”
G	Add in card PKG Break out
H	Add in card PKG

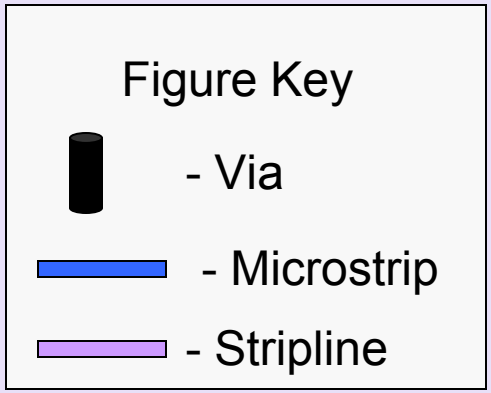

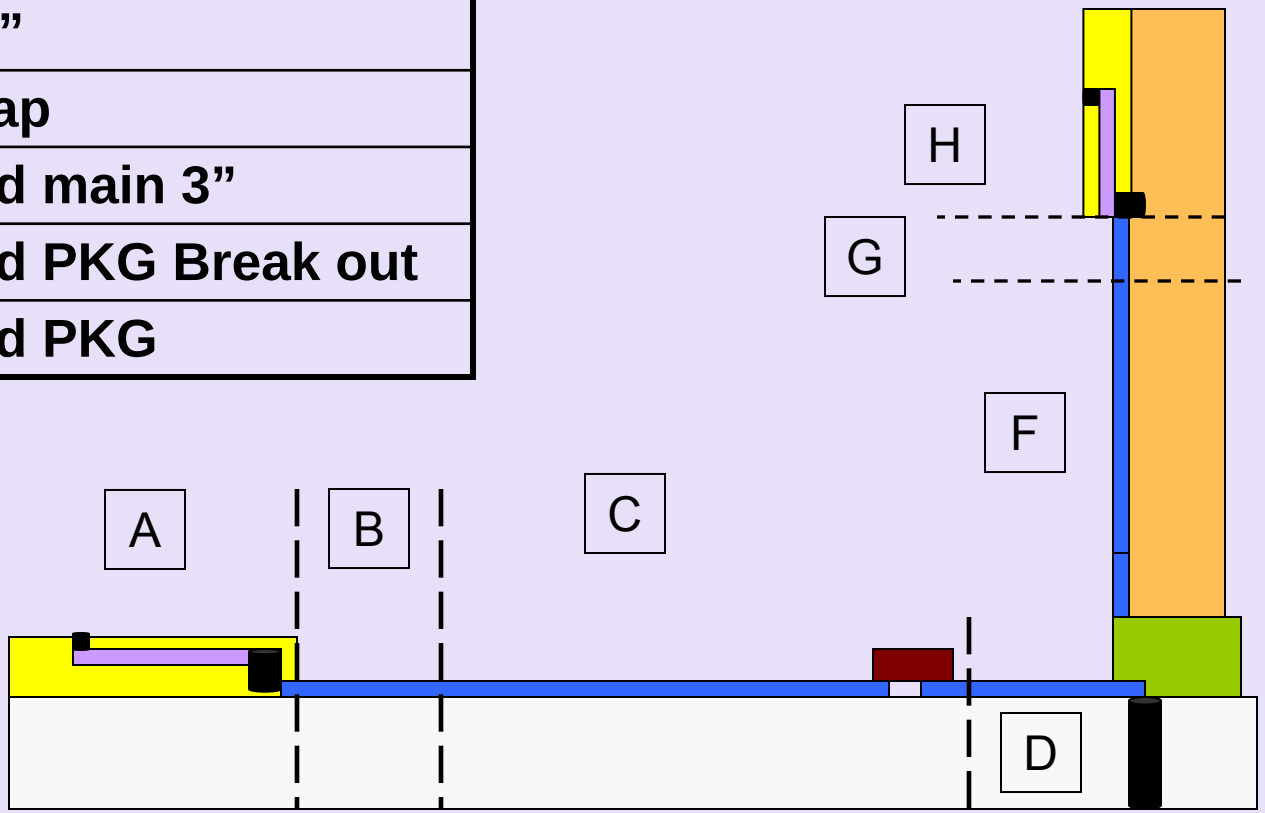

HVM Server Channel Configuration

- Two Connectors
- Mostly Stripline Routing
- 20” Total Trace Length
  - ✓ 4” AIC
  - ✓ 4” Riser
  - ✓ 16” Main Board

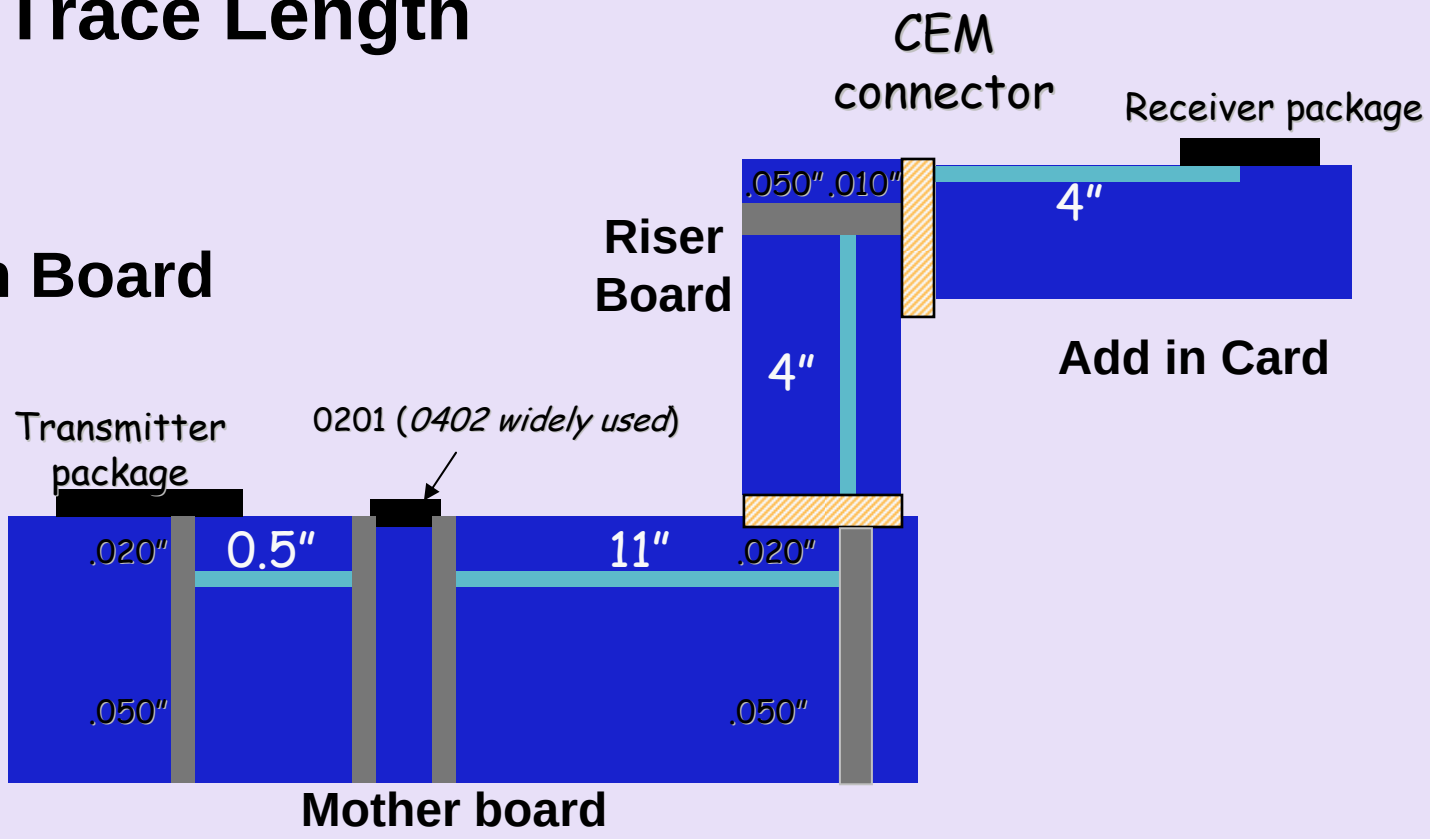

Client Channel - Frequency and Pulse Responses

- The insertion loss at 10GT/s is 6dB more than at 8GT/s
  - ✓ IL at 4GHz is -13.5dB (8GT/s)
  - ✓ IL at 5GHz is -19.3dB (10GT/s)

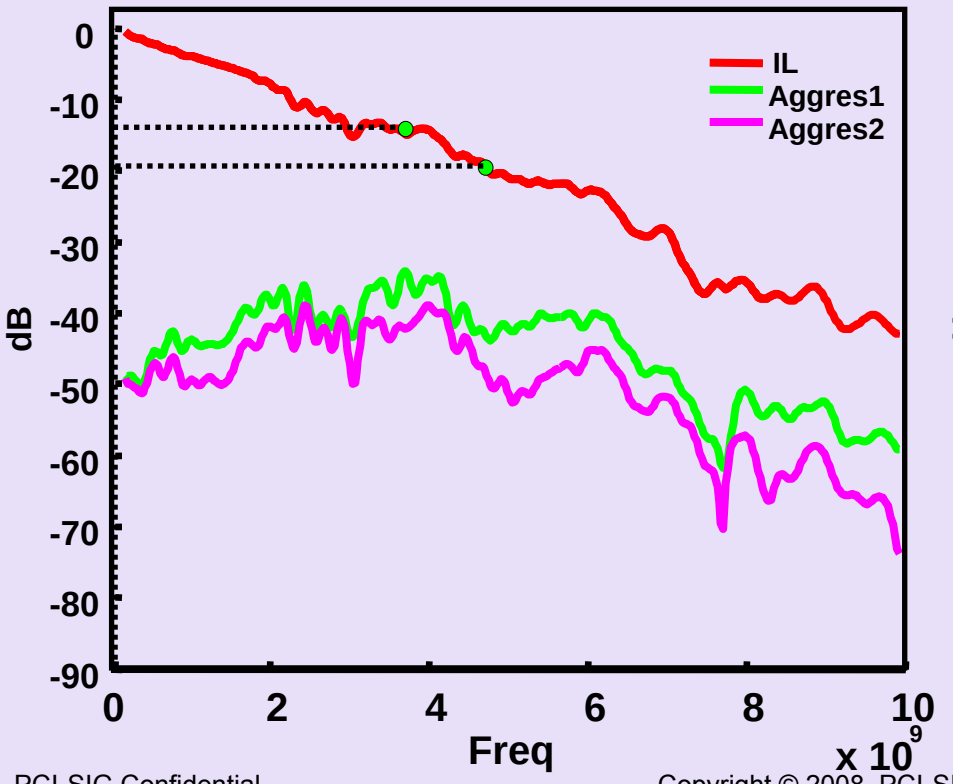

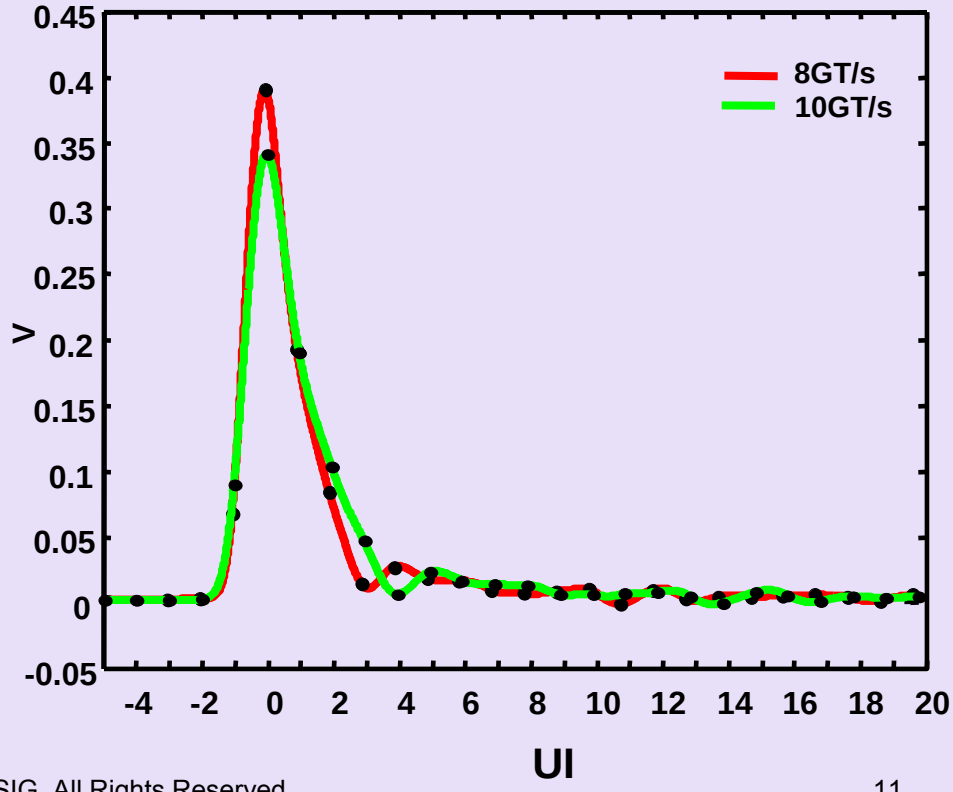

Sample BER Eye Diagrams

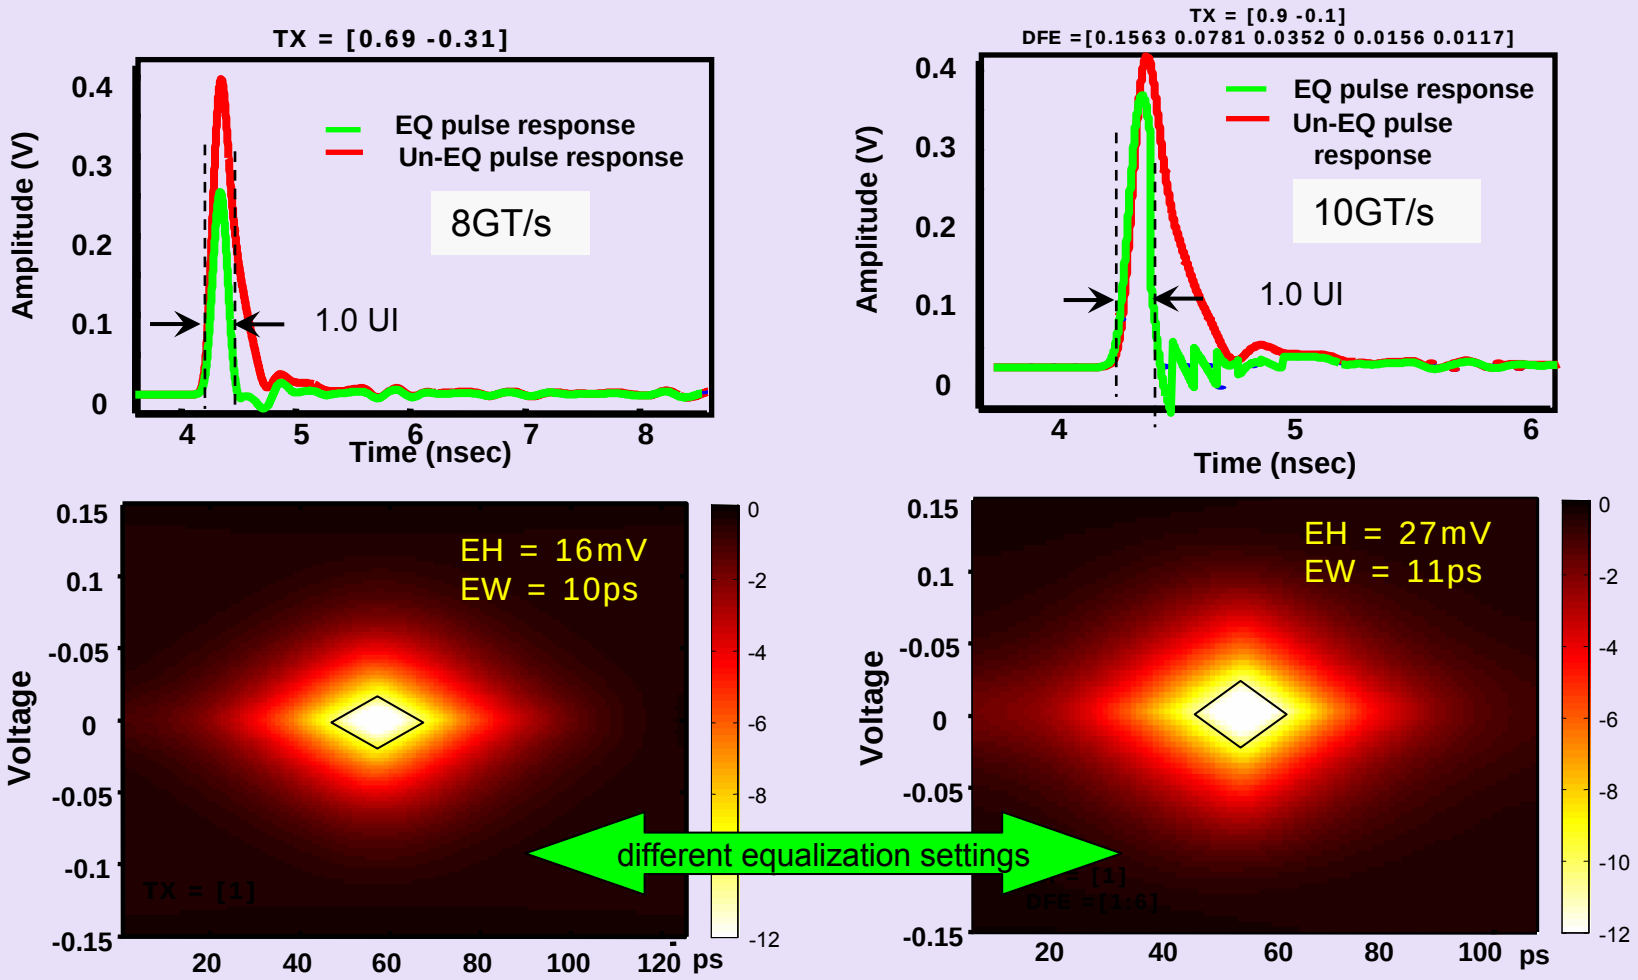

HVM Server Channel - Frequency and Pulse Responses

- ✓ IL at 4GHz is -16.5dB (8GT/s)
- ✓ IL at 5GHz is -18.4dB (10GT/s)

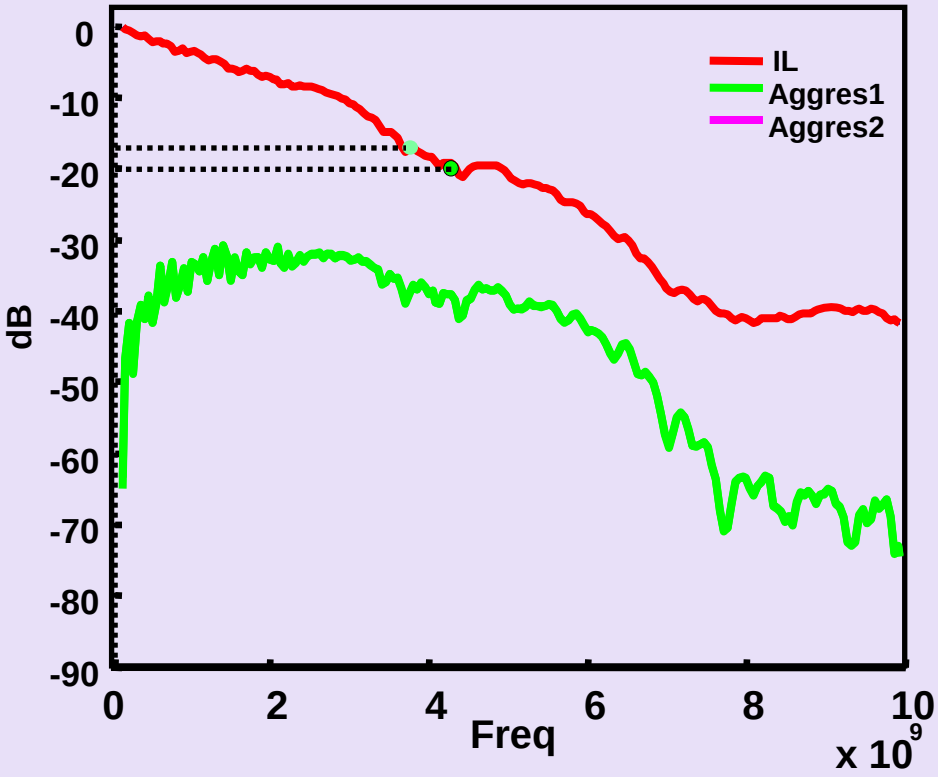

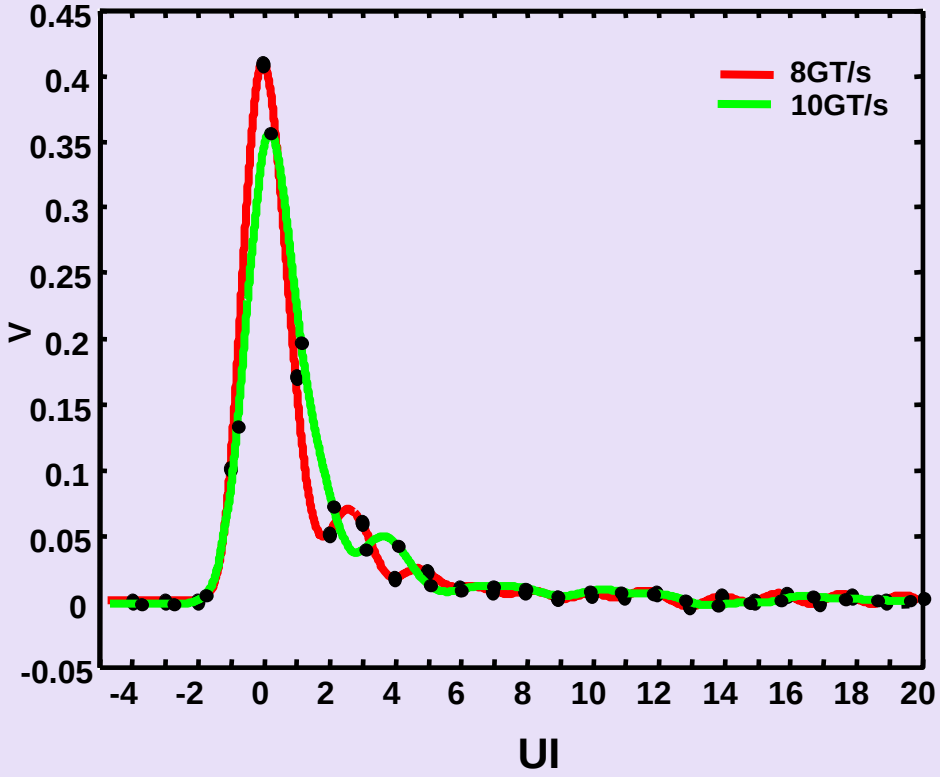

Simulation Results (Nominal)

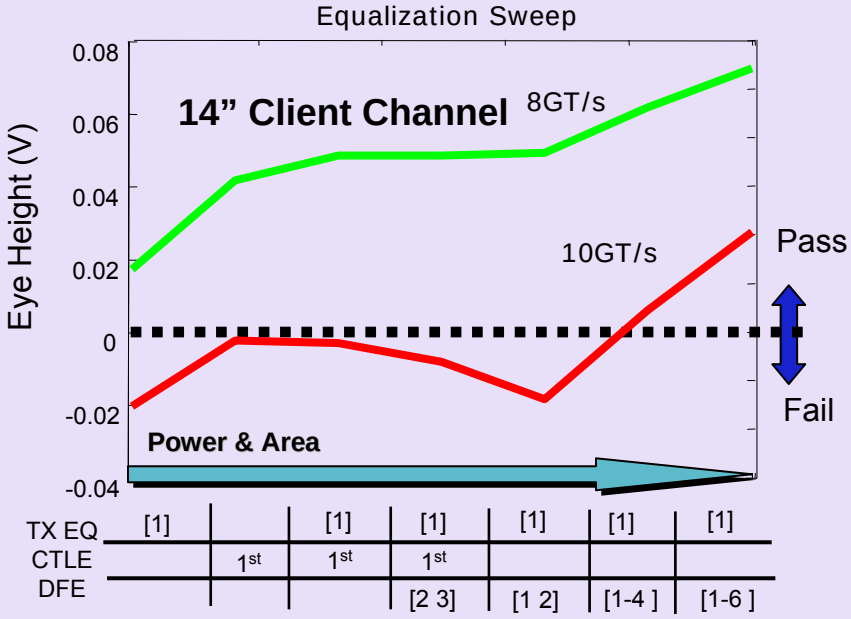

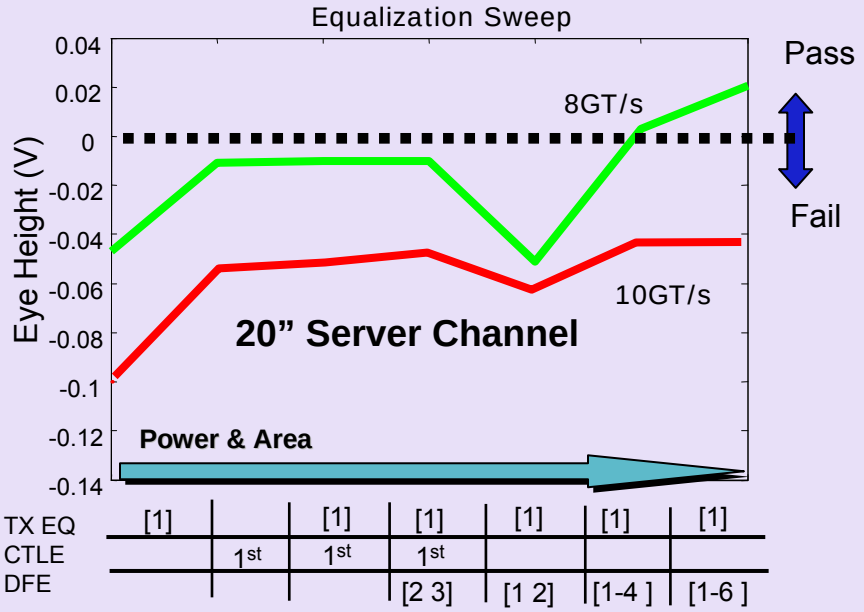

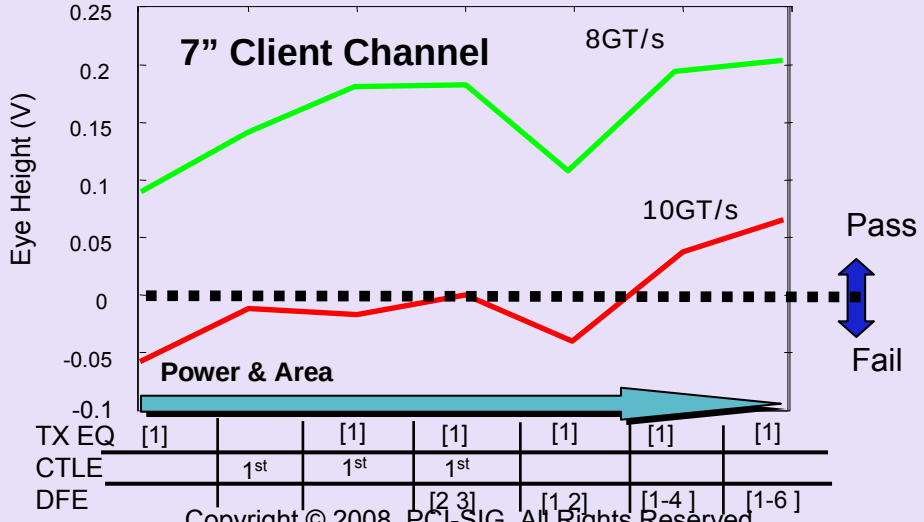

Simulation Results (Est W/C)

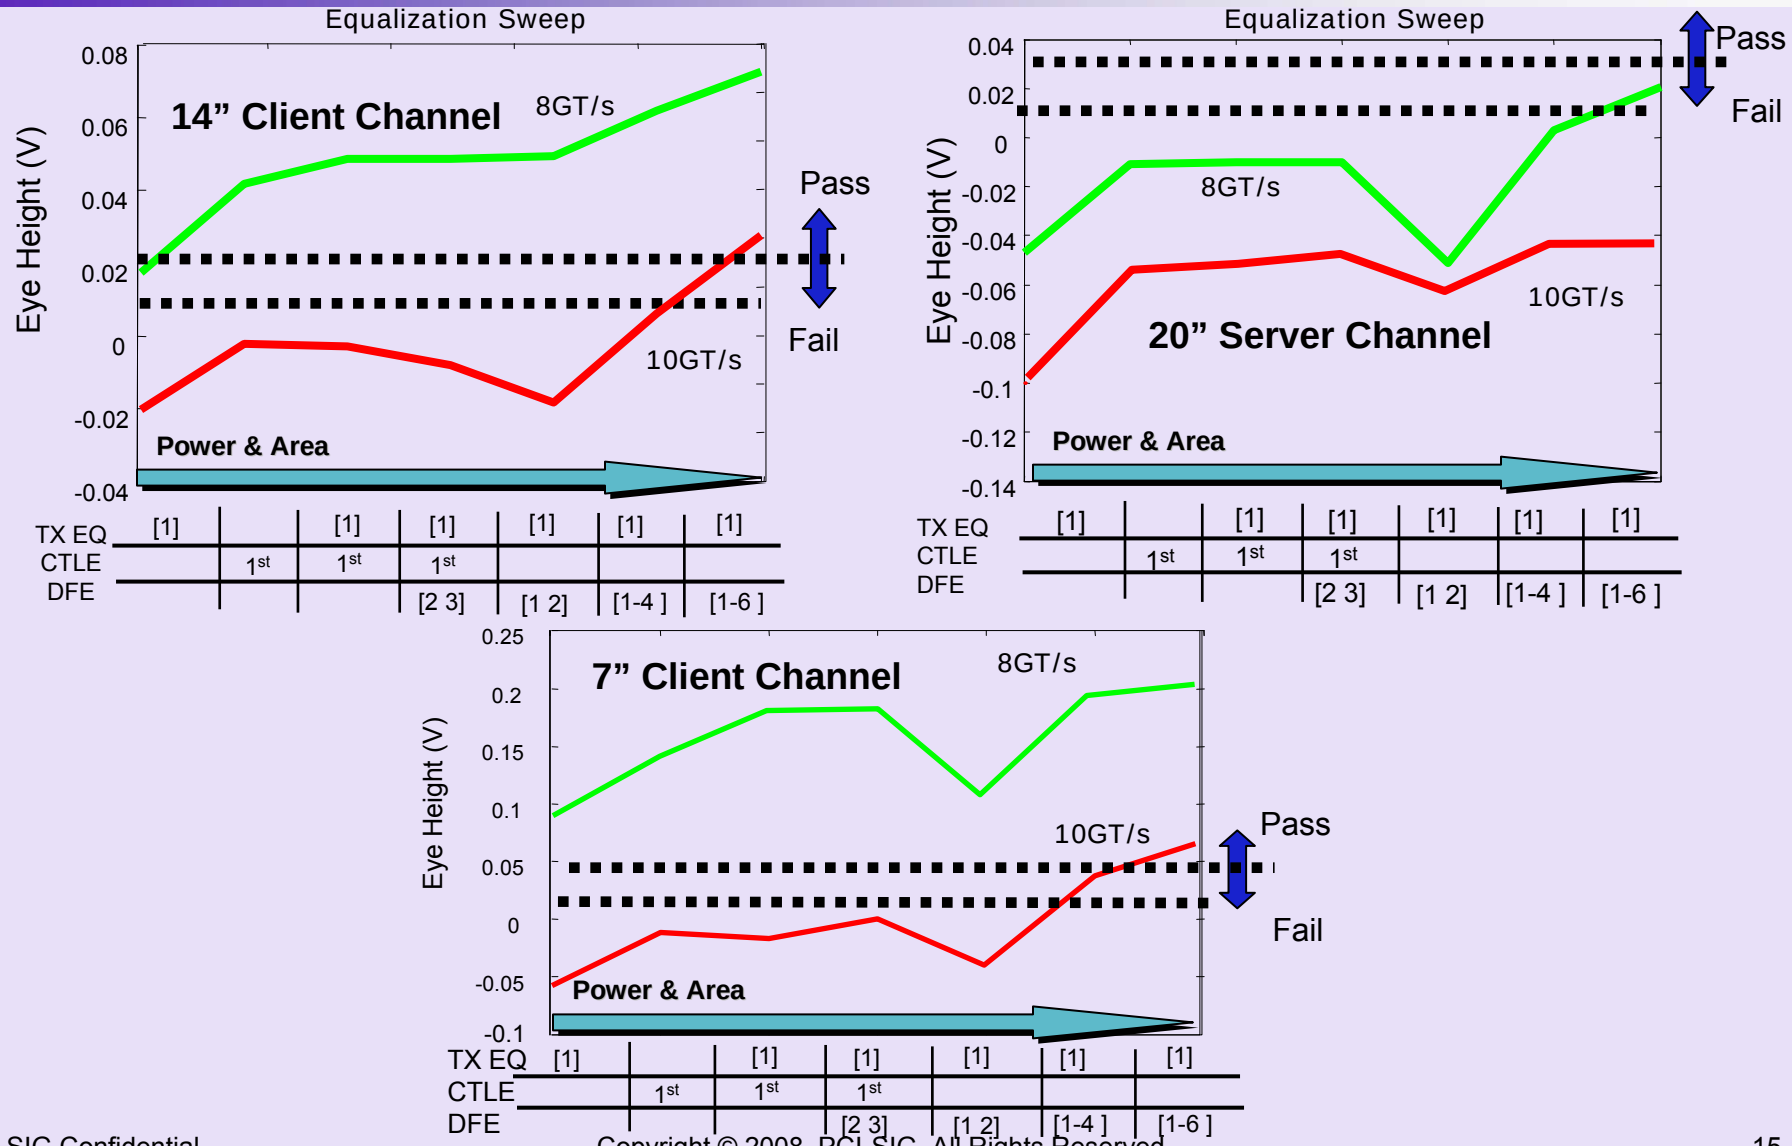

Rate Selection Summary

- 8GT/s is feasible over channels of interest with reasonable equalization
- 10GT/s imposes a power penalty
  - ✓ 8G-10G power increase somewhere between linear and quadratic
- 10GT/s imposes a cost penalty
  - ✓ Lower loss PCB materials
  - ✓ Backdrilled vias
  - ✓ Layout restrictions

PCIe 3.0 Tx Spec Subsection

- Transmitter Electrical parameters
  - ✓ Transmit PLL Characteristics
  - ✓ Tx Specification Location
  - ✓ Tx Timing Specifications
  - ✓ Adaptive TX Equalization?

Transmit PLL Characteristics

- 8.0 GT/s requires Tx PLL bandwidth and jitter peaking to be more tightly controlled than for 5.0 GT/s

- 2.0 Mhz 3dB Peaking
  - ✓ SSC limits low end
- 4.0 Mhz 3dB Peaking
- 5.0 Mhz 1dB Peaking
- 8.0 Mhz 3dB Peaking
- 16 Mhz 3dB Peaking

Base Spec TX Spec Location

- TX specification at silicon pins (2.0 base location)
  - ✓ Too difficult to quantify package interaction with unknown channel
- TX specification at die pad
  - ✓ Current spec direction
  - ✓ All relevant parameters can be specified at point that is independent of package and channel
  - ✓ Direct measurements not possible
    - – Standard de-embedding algorithm/methodology needed in base spec.
- TX specification at the end of reference channel(s)
  - ✓ Other option discussed in EWG
  - ✓ TX is compliant if it can produce passing signaling through a worst case channel(s)
  - ✓ Can a small number of reference channels capture all worst case Tx/package/channel interactions?
  - ✓ Contributions from various TX variables not clearly separated

Transmitter specs

Parameter	Description	5.0 GT/s	8.0 GT/s
UI	unit interval	200 ps ±300 ppm	125 ps ±300 ppm
$V_{TX-DIFF-PP}$	Differential p-p Voltage Swing	.8 – 1.2 V (pins)	.1 – 1.2 V (die)
$V_{TX-RESOLUTION}$	Minimum Resolution For Voltage Adjustments	N/A	50 mV
$T_{TX-1UI-RJ-8G}$	Rj over 1UI Width	N/A	.48 ps RMS max
$T_{TX-2UI-RJ-8G}$	Rj over 2UI Width	N/A	TBD
$T_{TX-UI-DJ-8G}$	Per UI Deterministic Jitter (1.5 Ghz +)	N/A	4 ps max
$T_{TX-HF-RJ-8G}$	TX Random Jitter (10 Mhz – 1.5 Ghz)	1.4 ps RMS max	1.6 ps RMS max
$T_{TX-HF-DJ-DD-8G}$	HF TX Deterministic Jitter	30 ps max	7 ps max
$T_{TX-LF-RMS-8G}$	LF TX Jitter (10 Khz – 10 Mhz)	3.0 ps RMS max	TBD

Substantial differences between 5.0 and 8.0 GT/s based on need to account for additional jitter effects (jitter amplification, etc)

Transmitter specs continued

Parameter	Description	5.0 GT/s	8.0 GT/s
$PKG_{TX-DIE-CAP}$	Equivalent Package Die Capacitance	N/A	1 pf Max
$PKG_{TX-PIN-CAP}$	Equivalent Package Pin Capacitance	N/A	.5 pf Max
PKG_{TX-LEN}	Equivalent Package Length	N/A	50 – 1500 mils
$Z_{TX-DIFF-DC}$	DC differential TX Impedance	N/A	120 ohm max
$L_{TX-SKEW}$	Lane-to-Lane Output Skew	500 ps + 4UI max	TBD
C_{TX}	AC Coupling Capacitance	75 – 200 nf	180 – 200 nf

- TX Equalization
  - ✓ 2 or 3 tap
  - ✓ Adjustable coefficients may be required
    - – Complicates TX silicon and form factor testing

Refclk Spec Subsection

- Reference Clock Electrical parameters
  - ✓ Refclk Architectures
  - ✓ Post processing steps
  - ✓ Jitter definitions

Clock Architectures

- PCIe Base spec defines two distinct Refclk architectures at 5.0 GT/s and 8.0 GT/s: common clock and data clocked
  - ✓ At 2.5 GT/s spec does not differentiate between 2 cases, but implicitly supports both
- Jitter margins for the two differ at 5.0 GT/s -- same at 8.0 GT/s.
  - ✓ PLL and CDR bandwidth changes remove any difference in jitter values between two architectures

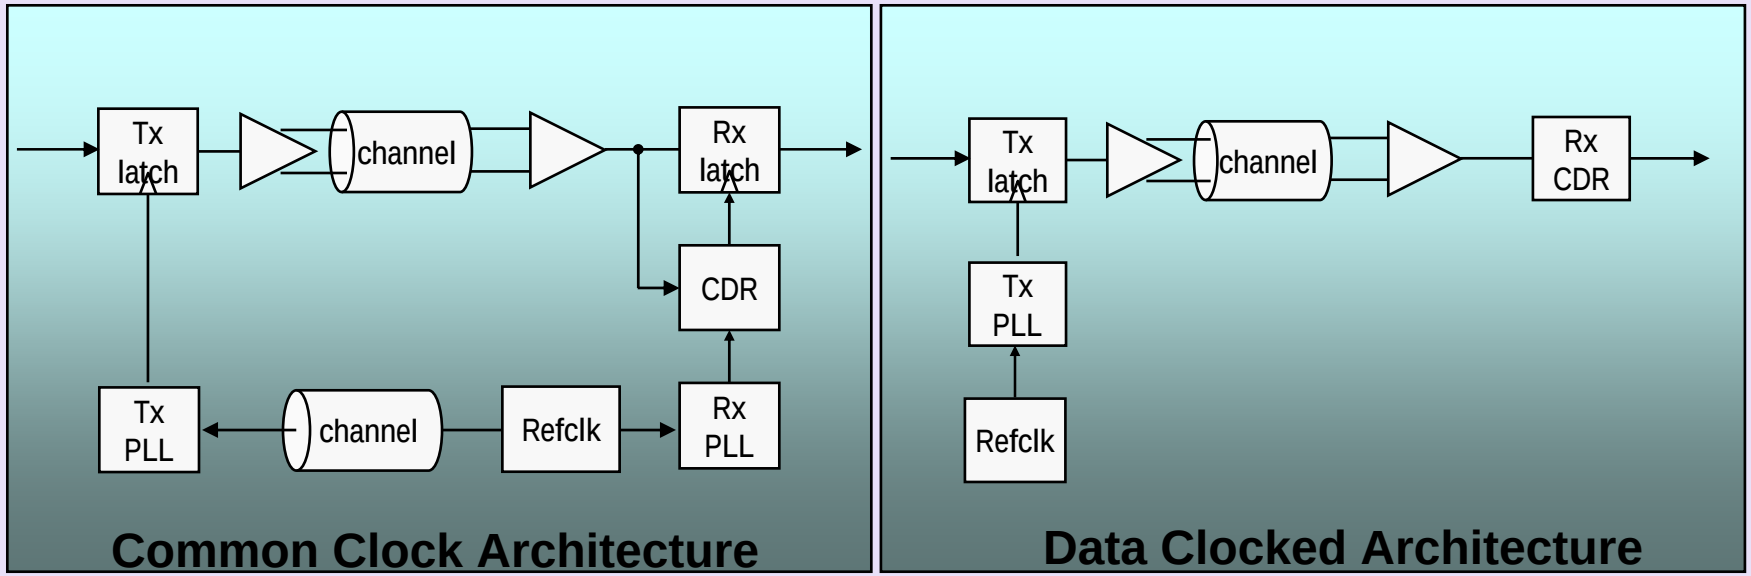

Refclk Post Processing for 8.0 GT/s

- Post processing removes jitter components that are measurement artifacts or otherwise irrelevant
- This process is NOT clock architecture dependent

	Common Clocked and Data Clock
< 10 MHz jitter components	No SSC removal PLL difference function (or min PLL) 0.01- 10 MHz step BPF
> 10 MHz jitter components	PLL difference function (or max PLL) 10 MHz step HPF Edge filtering

- PLL diff function: Difference between min and max PLL bandwidths
- Edge filtering: Smoothing function to reduce effects of sampling aperture inaccuracy
- Step filter Separates jitter into <10 MHz and ≥10 MHz bins

Reference Clock Data

- Obtained Connector Reference Clock Data With Several PCI Express 2.0 Systems
  - ✓ Measured with PCI-SIG® CLB 2.0 test fixture and RT scope.
- Analyzed HF Jitter with PCIe 2.0 and 3.0 Filters
  - ✓ 2.0 (3.1 ps RMS limit)
    - H1 16 Mhz, 3db Peaking, 40 db/dec rolloff
    - H2 5 Mhz, 1db Peaking, 40 db/dec rolloff
    - H3 1.5 Mhz High Pass Step.
  - ✓ 3.0 (1.0 ps RMS limit)
    - H1 4 Mhz, 3db Peaking, 40 db/dec rolloff
    - H2 2 Mhz, 3db Peaking, 40 db/dec rolloff
    - H3 10 Mhz Step

System A

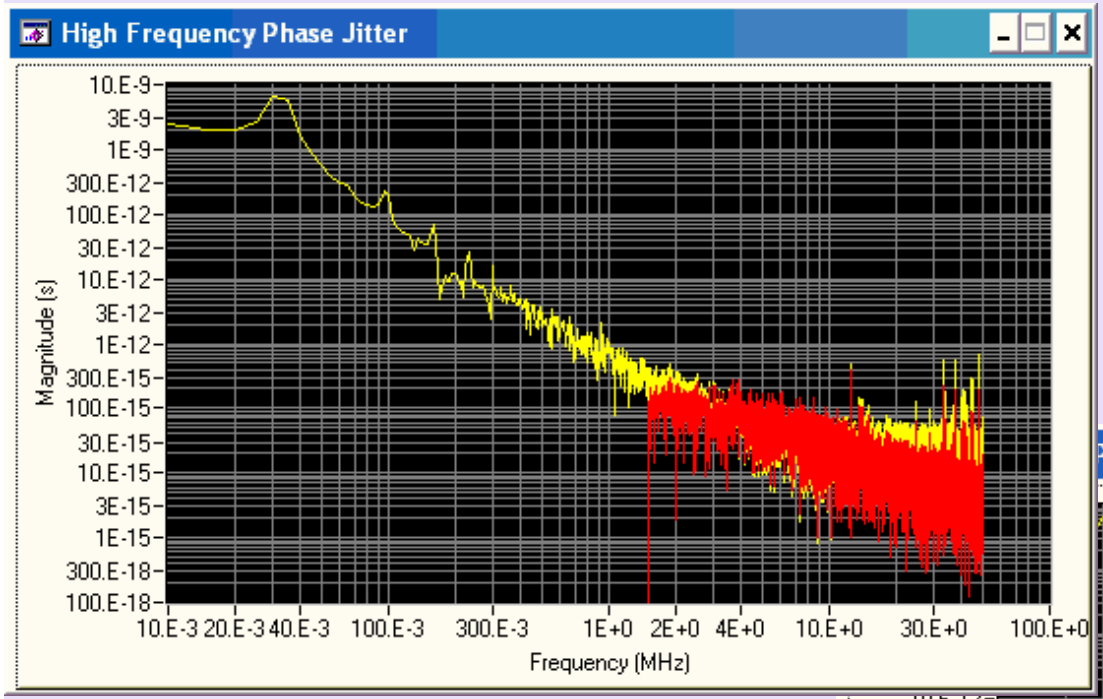

- PCIe 2.0 filter
  - ✓ 2.22 ps RMS

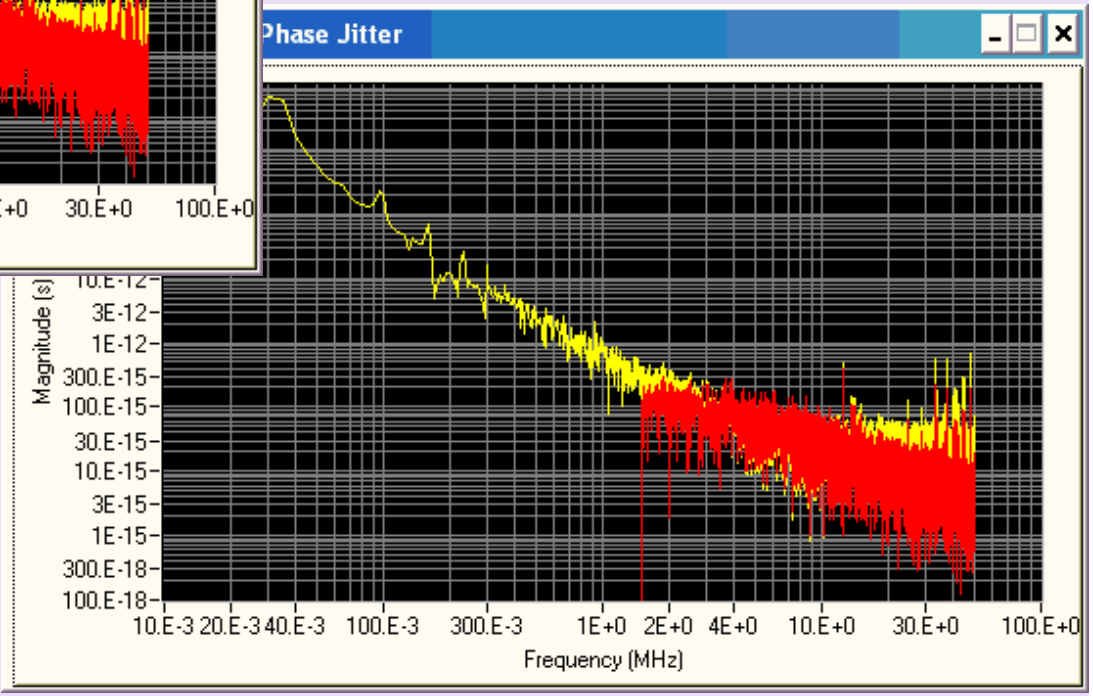

- PCIe 3.0 filter
  - ✓ .24 ps RMS
  - ✓ Existing compliant PCIe 2.0 systems can meet 3.0 HF jitter limits

System B

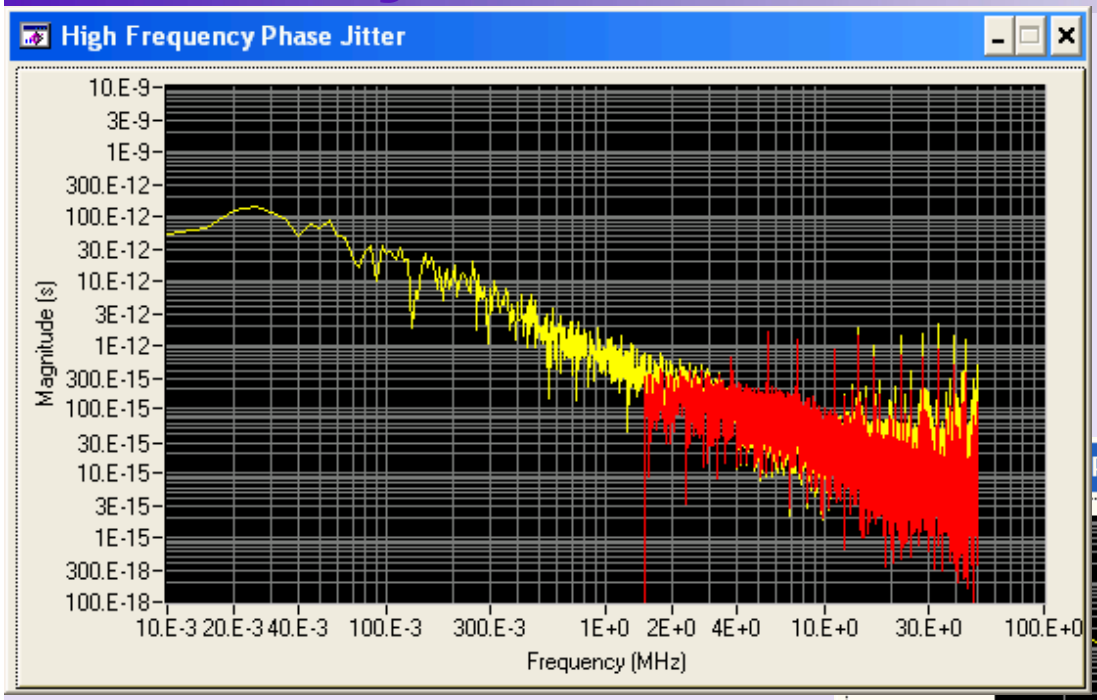

- PCIe 2.0 filter
  - ✓ 4.21 ps RMS

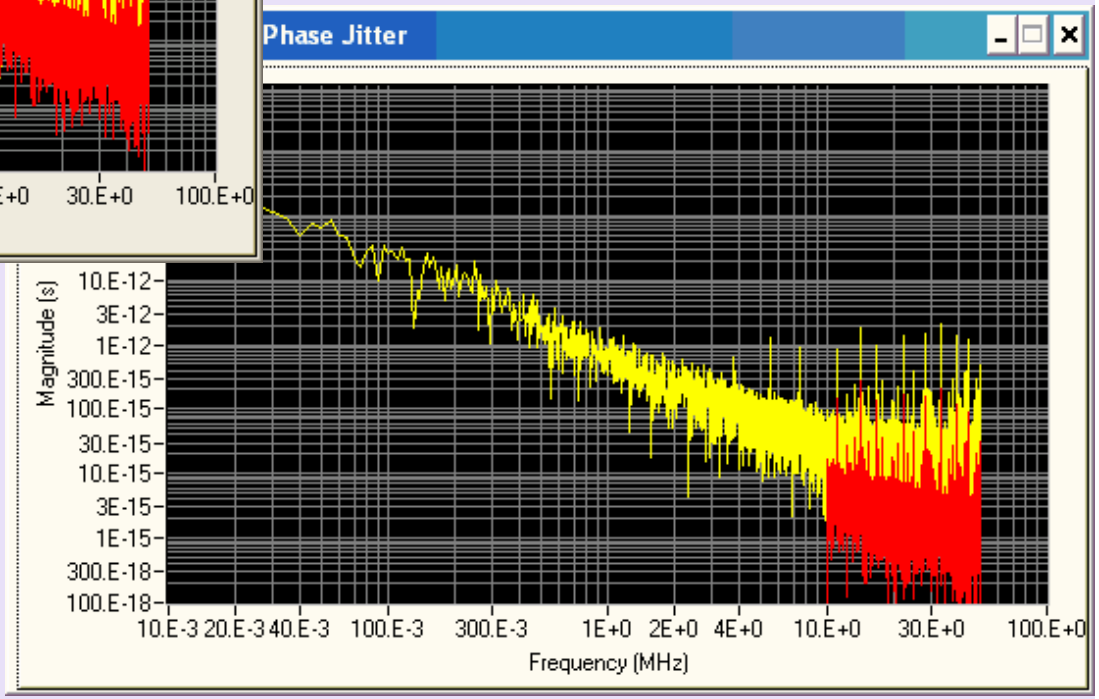

- PCIe 3.0 filter
  - ✓ .45 ps RMS
  - ✓ PCIe 2.0 HF limits are often more restrictive than 3.0 limits

System C

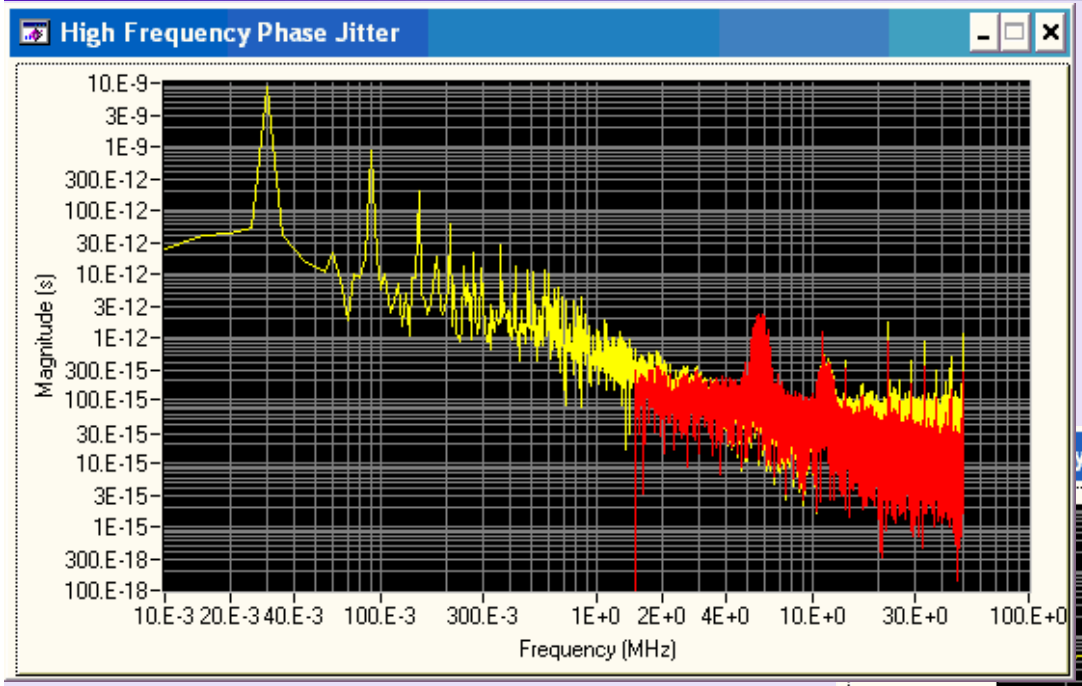

- PCIe 2.0 filter
  - ✓7.25 ps RMS

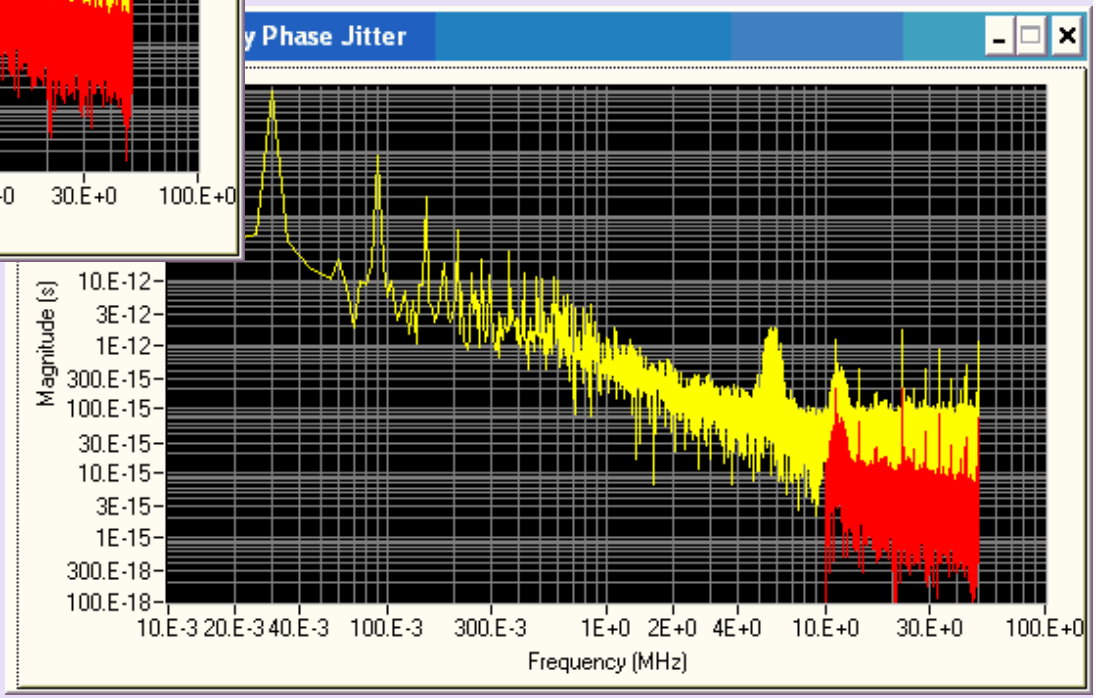

- PCIe 3.0 filter
  - ✓1.25 ps RMS

PCIe 3.0 Channel Spec – Major Changes

- Tx package defined in terms of C_{DIE}, C_{PAD}, and a swept length
- Rx package defined in terms of C_{DIE}, C_{PAD}, and a swept length
- Tx jitter is defined in terms of Dj and an Rj distribution
- Statistical simulation tools used to capture TX, channel, RX interactions
- A reference Rx equalization algorithm is applied to raw data as it appears at the Rx die pad

PCIe 3.0 Rx Spec Subsection

- PCIe 3.0 Receiver Specification
  - ✓ Major Change Summary
  - ✓ Scrambling Impact
  - ✓ RX Measurement Methodology

Major RX Specification Changes

- Jitter and voltage limits referenced to die pad
- Rx PLL bandwidth reduced to 2-4 Mhz.
- RX CDR bandwidth increased to 10 Mhz minimum.
- Jitter defined with bandlimited TJ and Dj components
- RX return loss replaced with C_{DIE}, C_{PIN}, C_{LENGTH}
- Jitter measured after applying inverse equalization algorithm

Base Spec Rx Equalization

- RX equalization is required.
- A specific RX equalization algorithm/method is not required by the specification.
- It is expected that most designs will be able to pass receiver base spec requirements with a simple technique like single pole CTLE.
- Impact on RX Measurement Methodology (Tolerance Test)
  - ✓ Apply baseline receiver equalization algorithm to calibrate test source OR
  - ✓ Calibrate noise sources with open eye and assume linearity as sources are increased
- Impact on form factor specifications
  - ✓ May have to apply baseline receiver equalization algorithm as part of TX data post processing.

Impact of Scrambling

- PHY Impact
  - ✓ Statistical DC balance only: DC wander
  - ✓ Statistical transition density: CDR tracking
  - ✓ Both appear to be solvable with minor circuit changes

- Ongoing PHY Work
  - ✓ Determine magnitude of DC wander and potential need for mitigation in Tx or Rx
  - ✓ Quantify frequency wander for DD architecture in presence of SSC and no data edges

What is Baseline Wander?

- In an AC coupled data transmission system, low freq signal components are removed by the HPF
- The average or DC value of the signal becomes data pattern dependent
- This causes a ‘wandering’ average
- The severity of baseline wander is dependent on the cut-off freq of the HPF and the PSD of the signal below this cut-off

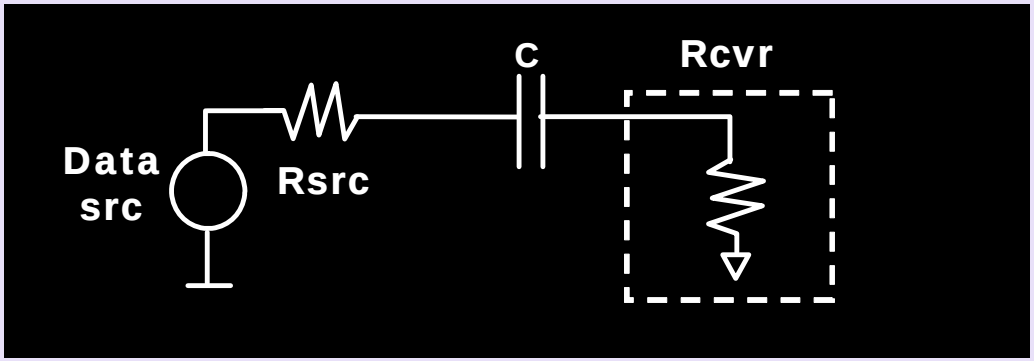

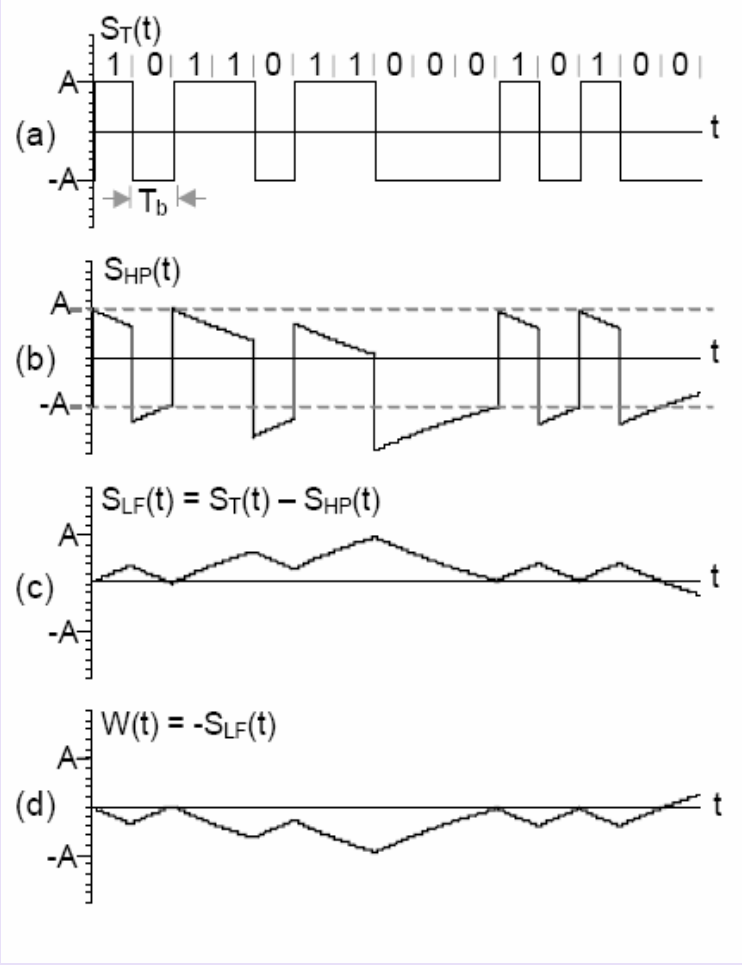

Simple Channel Model: With On-Die Capacitance

- 3 different HPF bandwidths
  - Case 1: A nominal capacitance 1pF with 100kW resistor for low cutoff
  - Case 2: A stretch (500kW) resistor case
  - Case 3: Similar to Case 1 with a 200nF AC line cap
- Sim conditions: 1.0 Vpp @ Tx, 10^6 random bits

Case #	R1 (Ω)	C1 (nF)	R2 (Ω)	C2 (pF)	R3 (KΩ)	R2-C1 BW (KHz)	R3-C2 BW (KHz)	BLW p-p (mV)
1	50	75	50	1	100	42.4	1591.6	112.5
2	50	75	50	2	500	42.4	159.2	33.5
3	60	200	60	1	100	13.3	1591.6	95

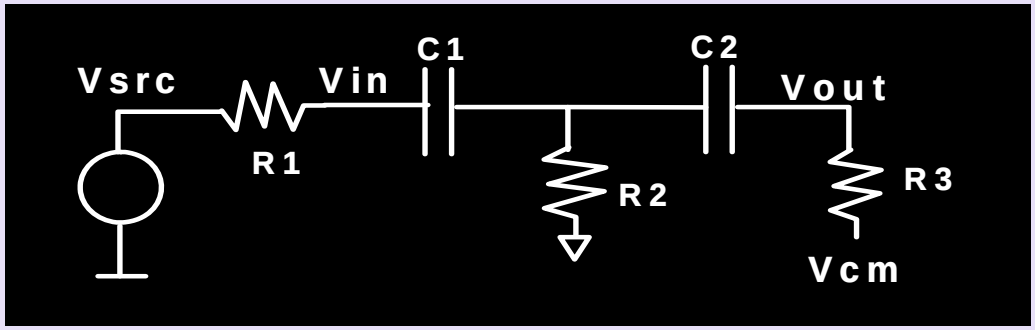

R1: source resistance

C1: off-chip capacitor

R2: termination resistance

C2: on-die capacitance

R3: on-die resistance

On Die RC Dominates Wander If On Die Capacitance Present

Baseline wander vs. On-Die HPF bandwidth

- Sweep on-chip RC keeping off-chip RC constant (R1=50Ω, R2=50 Ω, C1=75nF)
- As on-die HPF cut-off freq approaches off-chip bandwidth (=42.4 KHz), baseline wander reduction saturates as expected

R3 (Kohm)	C2 (pF)	R3-C2 BW (KHz)	BLW sigma (mV)	BLW p-p (mV)
1000	5	31.8	3.3	20
500	5	63.7	3.9	24.4
500	2	159.2	5	33.8
500	1	318.4	6.3	45
250	1	636.7	8.5	64.3
100	1	1591.8	13	106.1
50	1	3183.5	18.1	159.1
50	0.5	6367.1	25.3	243.1
10	1	15917.7	39.6	390.4
5	1	31835.4	55.8	504.7
5	0.5	63670.8	78.9	687.5

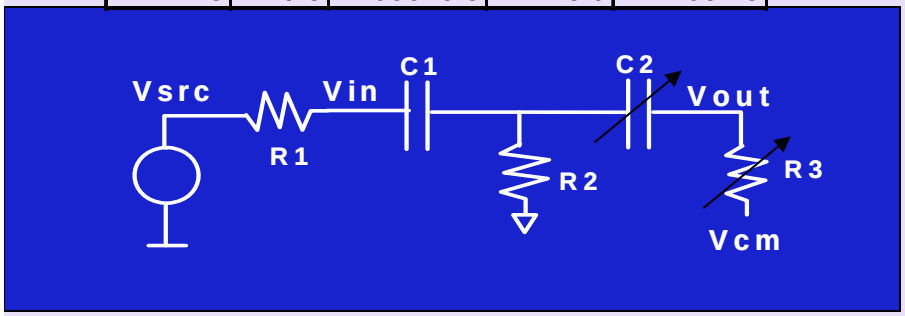

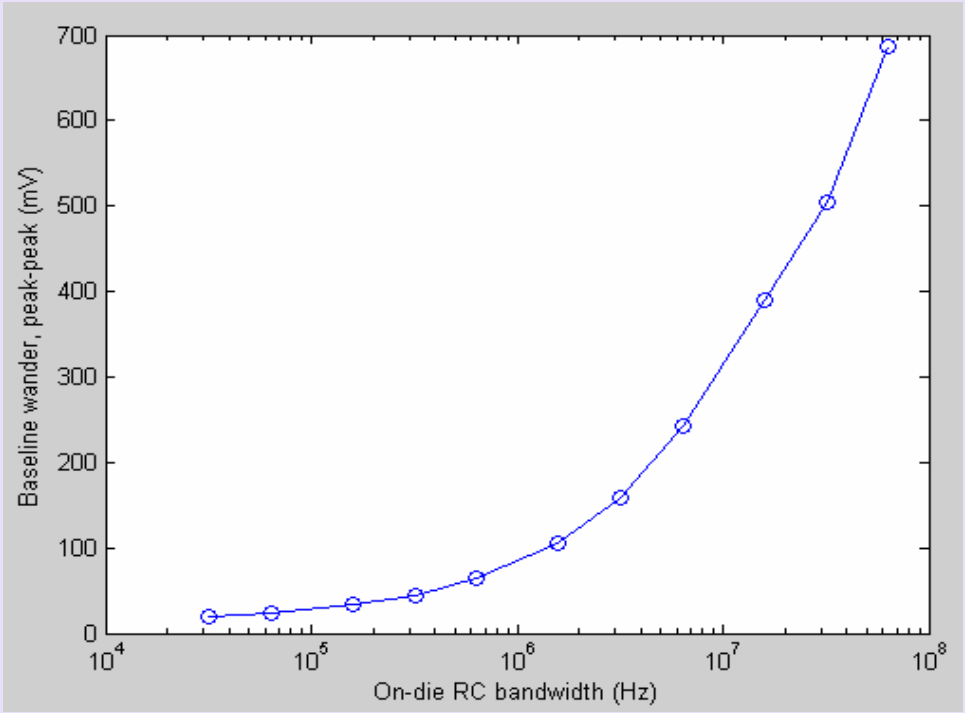

On Die RC Dominates Wander If On Die Capacitance Present

Effect of Transmit Equalization

- BLW scales linearly with transmit amplitude, i.e. it is a function of pre-aperture eye height
- Tx equalization attenuates low freq components resulting in reduced BLW
- Tx EQ sims:
  - 1 tap (postcursor) de-emphasis Tx Eq
  - Sweep tap coefficient for same Tx amplitude (1Vp-p)
  - BLW with and without on-chip cap are simulated (nominal case: R1=50 Ω, C1=75nF, R2=50Ω, C2=1pF, R3=100 Ω)

EQ setting	BLW p-p w/ on-chip cap (mV)	Pre-aperture eye height (V)	BLW p-p w/o on-chip cap (mV)
0	110	1.0	10.6
0.1	88	0.8	8.5
0.2	66	0.6	6.4
0.25	55	0.5	5.3
0.3	44	0.4	4.2
0.4	22	0.2	2.1

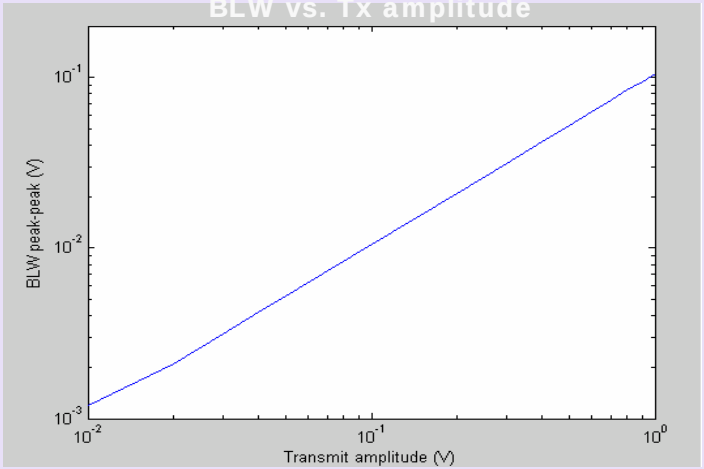

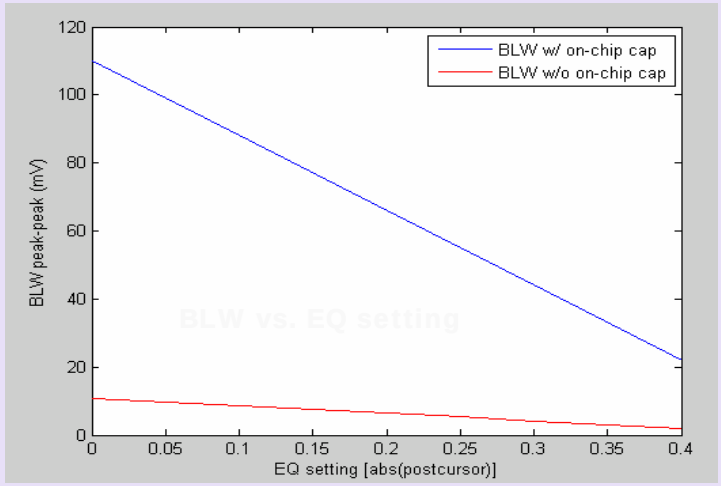

Base Line Wander Next Steps

- Ongoing simulation work to determine accurate worst case number.
- Analyze possible mitigation techniques
  - ✓ Bit Stuffing
  - ✓ DC restoration circuit in RX
  - ✓ DC coupled receiver
  - ✓ Combinations of above approaches
  - ✓ Other techniques?

Form Factor TX Measurement Methodology

- Option 1 - Specify standard fixture(s) requirements and include in determining form factor limits (CEM 2.0 methodology)
  - ✓ Pros
    - Don't need to specify de-embedding algorithm/procedure that can be applied consistently across industry
    - PCI-SIG can provide standard fixtures to members
  - ✓ Cons
    - Will require tight control of fixture parameters and likely add cost to fixtures
      - Fixtures may be high cost anyway if they have to provide receiver feedback to drive TX adaptive EQ to different states
      - Fixture cost still small relative to test equipment cost
    - May not be possible at 8 GT/s. (investigation needed)

- Option 2 – Specifying standard de-embedding process/requirements for any form factor fixture (don't include fixture in form factor limits)
  - ✓ Pros
    - A variety of fixtures with different characteristics could provide equivalent results
  - ✓ Cons
    - Need to specify de-embedding algorithm/procedure that can be applied consistently across industry
    - Getting accurate simulation results exactly at the edge finger/connector may be difficult

Form Factor Reference Clock Testing

- Option 1 – Test Reference Clock Separately
  - ✓ Pros
    - Simpler measurement setup than dual port
  - ✓ Cons
    - Removes ability to trade off clock and data jitter at system level
    - Must account for not having a clean reference clock for standard motherboard TX test
- Option 2 – Use Dual Port Simultaneously Clock/Data (Methodology Specified in CEM 2.0)
  - ✓ Pros
    - Allows tradeoff of data and clock jitter at system level
    - Don't have to worry about how to test real motherboard without clean clock
    - No issues testing with SSC on
  - ✓ Cons
    - More complex measurement setup – but already proven for CEM 2.0
    - Ability to trade off clock and data jitter adds little relief with clock jitter budget at 1 ps Rj (RSS with other other parts of RJ budget)

Form Factor Methodology For 3.0

- Need to investigate whether CEM 2.0 methodology for determining connector voltage/jitter limits will work for 3.0
  - ✓ Less margin available
- Additional constraints beyond jitter/voltage margin may be needed to preserve enough solution space for 3.0
  - ✓ TDR
  - ✓ Return Loss
  - ✓ Other . . .

Major Work Items Upcoming

- Demonstrate method of de-embedding to die pad
  - ✓ Good progress: several options being evaluated
- Close on Tx equalization choices
  - ✓ Trainable vs. fixed coefficients
- Resolve DC wander effects
  - ✓ Rx voltage margin, effective CDR BW impact
- Long server channel mitigation costs/effectiveness

Future Plans

- Rev0.3
  - ✓ Data rate, encoding set
  - ✓ Tx, Rx parameter tables
  - ✓ Being reviewed by EWG now
- Rev0.5
  - ✓ Tx, Rx reference planes defined
  - ✓ All parameters defined
  - ✓ Tx, Rx equalization defined
- Rev0.7
  - ✓ All parameter values stable
  - ✓ Statistical scripts included in spec
- Rev0.9
  - ✓ Minor formatting/typo edits

Backup

CEM 2.0 Methodology Review

System Board TX Eye Methodology

- Simulate end to end eye diagrams

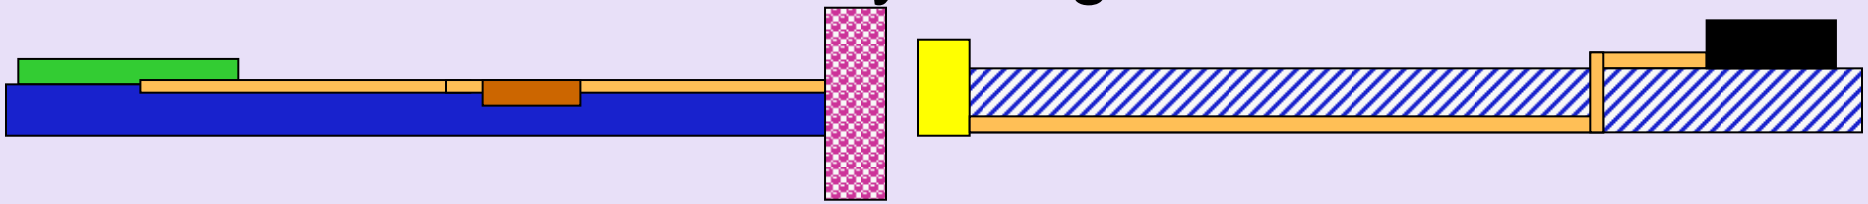

- Identify all end to end failures (worst case pattern)
  - 120 mVolt Eye Height (Base Spec Rx Pin Limit)
  - 142 ps Eye Width (Interconnect only) (Base Spec Channel Limit)

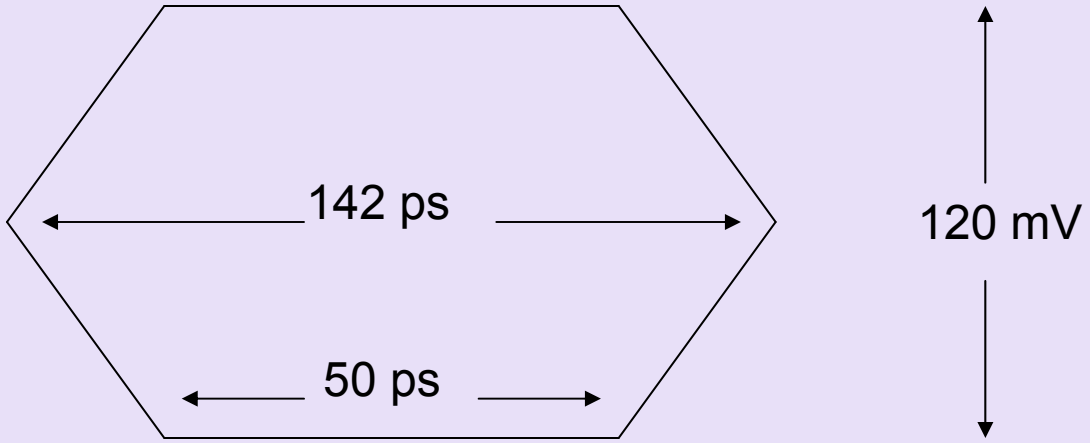

Worst Case Patterns

- Peak Distortion Analysis
  - ✓ Deterministically Calculates Worst Case Patterns Given
    - Channel S Parameters
    - Pulse Response
  - ✓ Used For Simulation Data In This Presentation
- Differences From Pseudo Random or CMM Patterns Can Be Very Large (~ 30 ps eye width)

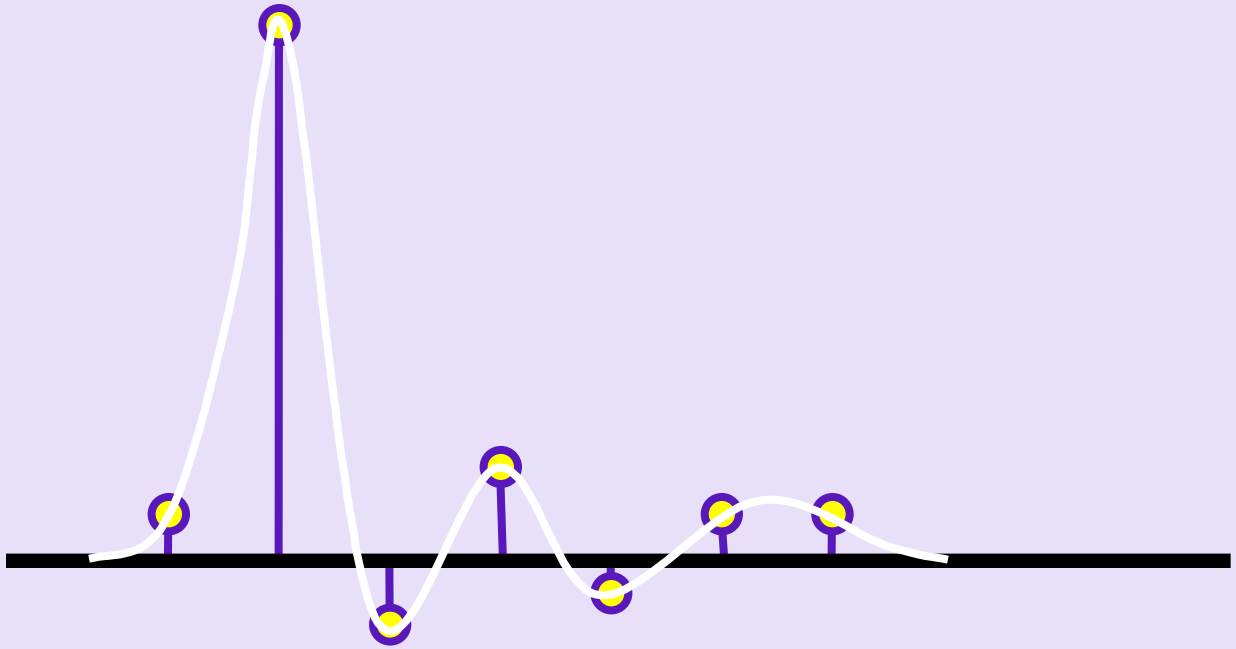

System Board TX Eye Methodology

- Simulate end to connector eye diagrams

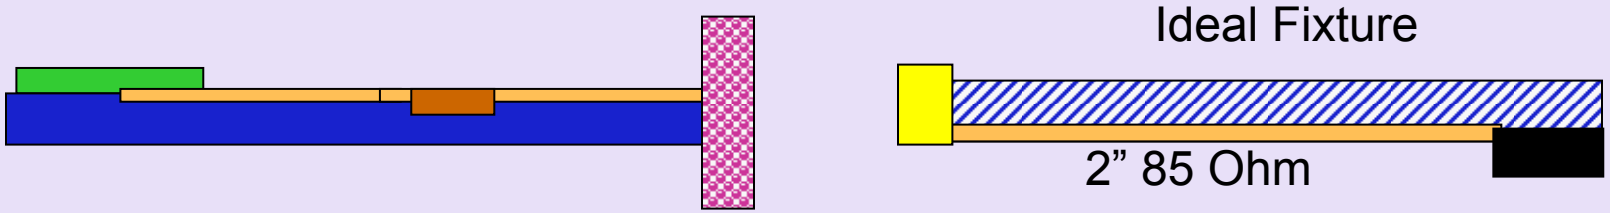

- Use CMM pattern as with real world test
- Correlate with end to end worst case pattern failures

- CEM eye specifications include ideal fixture
  - ✓ No need to de-embed if similar fixture used

Simulation Methodology

- The resultant eyes of the End to End and CEM simulations are plotted against each other for a large number of cases
- A Horizontal line is drawn with respect to the End to End eye to signify insufficient opening in the system
- A Vertical line is drawn such that no End to end failures are to the right
- Instances in the lower right quadrant would indicate End to End failures not screened out by CEM
- Instances in the upper left quadrant are cases which work End to End, but are screened out by the CEM*