

T10/08-404r2

SAS-2: Tools for TX characterization (and jitter tolerance setup qualification)

Mathieu Gagnon

Guillaume Fortin

PMC-Sierra



- Link to Previous Material
- Guiding Principles
- Transmitter Specifications
- Far-end Measurements Orthogonality
- Issues with SASWDP
- SAS_EYEOPENING
 - Presentation
 - Features
 - Operational Features
- SASWDP vs. SAS_EYEOPENING
- Summary
- To do

Link to Previous material (1)

- In 08-330r0 and 08-345r1 Kevin Witt presented results showing how the SASWDP program could be used to characterize transmitters at the far end, after the compliance channel.
 - The idea is to replace traditional near-end measurements with a single far-end metric
 - Makes the process more flexible to different TX features
 - Makes the process pattern-independent
 - Same point as for RX characterization - consistent
- The measurement is taken at the TX – as close as possible to the chip pin and convolved with the reference channel.

Guiding Principles (1)

- 2 ways to view the TX compliance:
 1. Generate a waveform that an RX can recover
 2. Meet specific TX criteria that will ensure an RX can recover its output
- The WDP approach is of the first kind – verify if the RX can recover without checking the specific parameters directly
- To verify specific TX criteria is restrictive on the type of equalization that can be provided by a TX

Guiding Principles (2)

- Spec. contains only few mandatory patterns to be generated by the TX
 - CJTPAT (Table 61, note g)
 - D10.2 or D21.5 (Table 61, note f)
 - PHY_TEST_PATTERN is not mandatory

- Specification measurements are usually done at the near end
 - Easier to setup
 - At 6 Gbs, de-embedding results at compliance points is an issue
 - Measurement at the chip removes uncertainty: Treats board as part of the channel
 - Minimal impact of testing probe – likely no de-embedding

Guiding Principles (3)

- Most TX specifications can be tested easily at the near end
 - Mostly DDJ remains, which is mostly orthogonal to other specifications
- Next slides review the standard TX specs:
 - Whether they are orthogonal to DDJ from TX and channel
 - How they relate to system performance
 - Whether they are covered easily by near-end measurements

Transmitter Specifications

Specification	Best Pattern	Near-End Measurement	Simulated Far-End Measurement	Comments
RJ	D10.2 or D24.3	Easy	Difficult Requires long capture	
BUJ	PJ	CJTPAT	Easy	Difficult Requires long capture
	DCDJ	CJTPAT or D10.2	Easy	Possible
	SSCJ	CJTPAT or D24.3	Easiest (but not easy)	Difficult Requires long capture
	DDJ	CJTPAT	Constrains de-emphasis type	Eye opening after DFE
Amplitude	CJTPAT	Possible	Tap0 magnitude in CPR	Zero-Length Test Load must be de-embedded if used to capture tx signal
De-Emphasis	CJTPAT	Constrains de-emphasis type	Eye opening after DFE	
Rise/Fall Times	CJTPAT	Easy Relevant for EMI	Eye opening after DFE	

Far-end measurements orthogonality


- Specs affected by channel are best measured at far-end
 - TX amplitude
 - TX De-emphasis
 - TX DDJ
 - TX Rise-Fall time
- RX eye opening is really a combination of all the specs.
 - Need a tool for far-end measurements that can ignore the effects of specs measured at the near-end.
 - I.e., the far-end specs must be measured orthogonally to the near-end ones.
- Can we use SASWDP?

- Requires a PRBS pattern (new to spec)
 - Data must be properly captured to be periodic
- Requires the digital pattern
 - Data and pattern must be aligned
- Has large variations vs. pattern
 - Up to 0.8 dB between two PRBS10
- Sometimes the Clock Recovery does not converge
- Its CDR creates a new « hidden » spec.
- ncDDJ incorrectly measured?
 - Large variations anyway
- See next slide

Issues with SASWDP

TX	TX de-emphasis	Pattern	Channel	SASWDP (original)	ncDDJ (original)
Generated 113ps rft 0-100%	3	<i>TxDatFile_6m-cjtpat</i>	10m SAS	11.4	0.578
Generated 113ps rft 0-100%	3	<i>TxDatFile_10m-prbs10</i>	10m SAS	10.6	0.297
Generated 113ps rft 0-100%	3	PRBS10, poly x204	10m SAS	11.2	0.263
Generated 113ps rft 0-100%	0	<i>TxDatFile_6m-cjtpat</i>	10m SAS	11.3	0.52
Generated 113ps rft 0-100%	0	<i>TxDatFile_10m-prbs10</i>	10m SAS	13.2	0.325
Generated 113ps rft 0-100%	0	PRBS10, poly x204	10m SAS	14	0.364
Generated 113ps rft 0-100%	3	<i>TxDatFile_6m-cjtpat</i>	HP24	15.4	0.725
Generated 113ps rft 0-100%	3	<i>TxDatFile_10m-prbs10</i>	HP24	10.6	0.368
Generated 113ps rft 0-100%	3	PRBS10, poly x204	HP24	10.3	0.397
Generated 113ps rft 0-100%	0	<i>TxDatFile_6m-cjtpat</i>	HP24	24.3	1
Generated 113ps rft 0-100%	0	<i>TxDatFile_10m-prbs10</i>	HP24	12	0.284
Generated 113ps rft 0-100%	0	PRBS10, poly x204	HP24	11.7	0.324

- SAS_EYEOPENING.m script developed
- Evaluates pulse response from the channel (-4 precursors to +40 post-cursors – inspired by SASWDP)
- Extracts sampling instant assuming a random input
 - Insensitive to input pattern
 - Could be improved to consider 8b10b?
- Re-computes the pulse response at this point
- Computes a simple « worst-case » 8b10b sequence
- Computes the eye opening due to DDJ, after a perfect 3-taps DFE
- Outputs information about each of the DFE's 3-taps compensation




SAS_EYEOPENING vs. SASWDP

TX	TX de-emphasis	Pattern	Channel	SASWD P (original)	ncDDJ (original)	"8b10b" eye opening
Generated 113ps rft 0-100%	3	<i>TxDataFile_6m-cjtpat</i>	10m SAS	11.4	0.578	65.6%
Generated 113ps rft 0-100%	3	<i>TxDataFile_10m-prbs10</i>	10m SAS	10.6	0.297	65.6%
Generated 113ps rft 0-100%	3	PRBS10, poly x204	10m SAS	11.2	0.263	65.5%
Generated 113ps rft 0-100%	0	<i>TxDataFile_6m-cjtpat</i>	10m SAS	11.3	0.52	56.7%
Generated 113ps rft 0-100%	0	<i>TxDataFile_10m-prbs10</i>	10m SAS	13.2	0.325	56.6%
Generated 113ps rft 0-100%	0	PRBS10, poly x204	10m SAS	14	0.364	56.6%
Generated 113ps rft 0-100%	3	<i>TxDataFile_6m-cjtpat</i>	HP24	15.4	0.725	65.4%
Generated 113ps rft 0-100%	3	<i>TxDataFile_10m-prbs10</i>	HP24	10.6	0.368	64.6%
Generated 113ps rft 0-100%	3	PRBS10, poly x204	HP24	10.3	0.397	64.6%
Generated 113ps rft 0-100%	0	<i>TxDataFile_6m-cjtpat</i>	HP24	24.3	1	61.4%
Generated 113ps rft 0-100%	0	<i>TxDataFile_10m-prbs10</i>	HP24	12	0.284	60.7%
Generated 113ps rft 0-100%	0	PRBS10, poly x204	HP24	11.7	0.324	60.6%

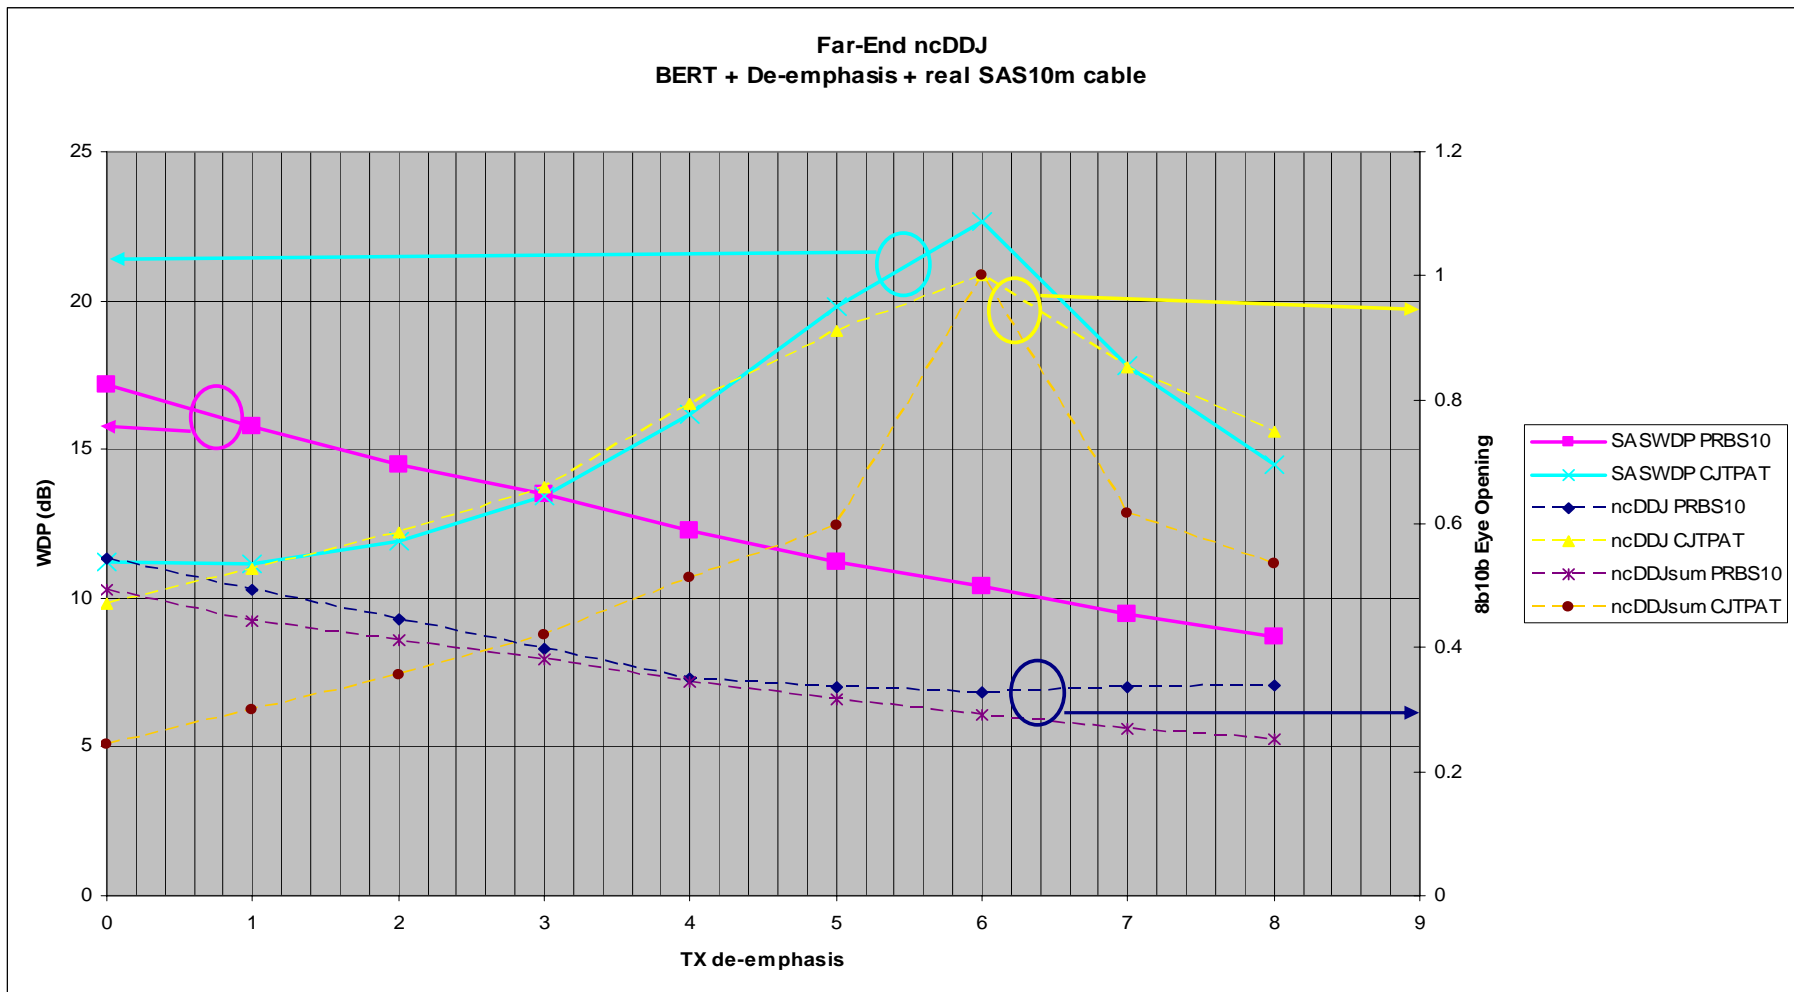
SAS_EYEOPENING vs. SASWDP

- SASWDP vs. SAS_EYEOPENING (PRBS10)




SAS_EYEOPENING vs. SASWDP

- SASWDP: CJTPAT vs. PRBS10 (results averaged)
 - PRBS10 gives consistent results vs. De-emphasis



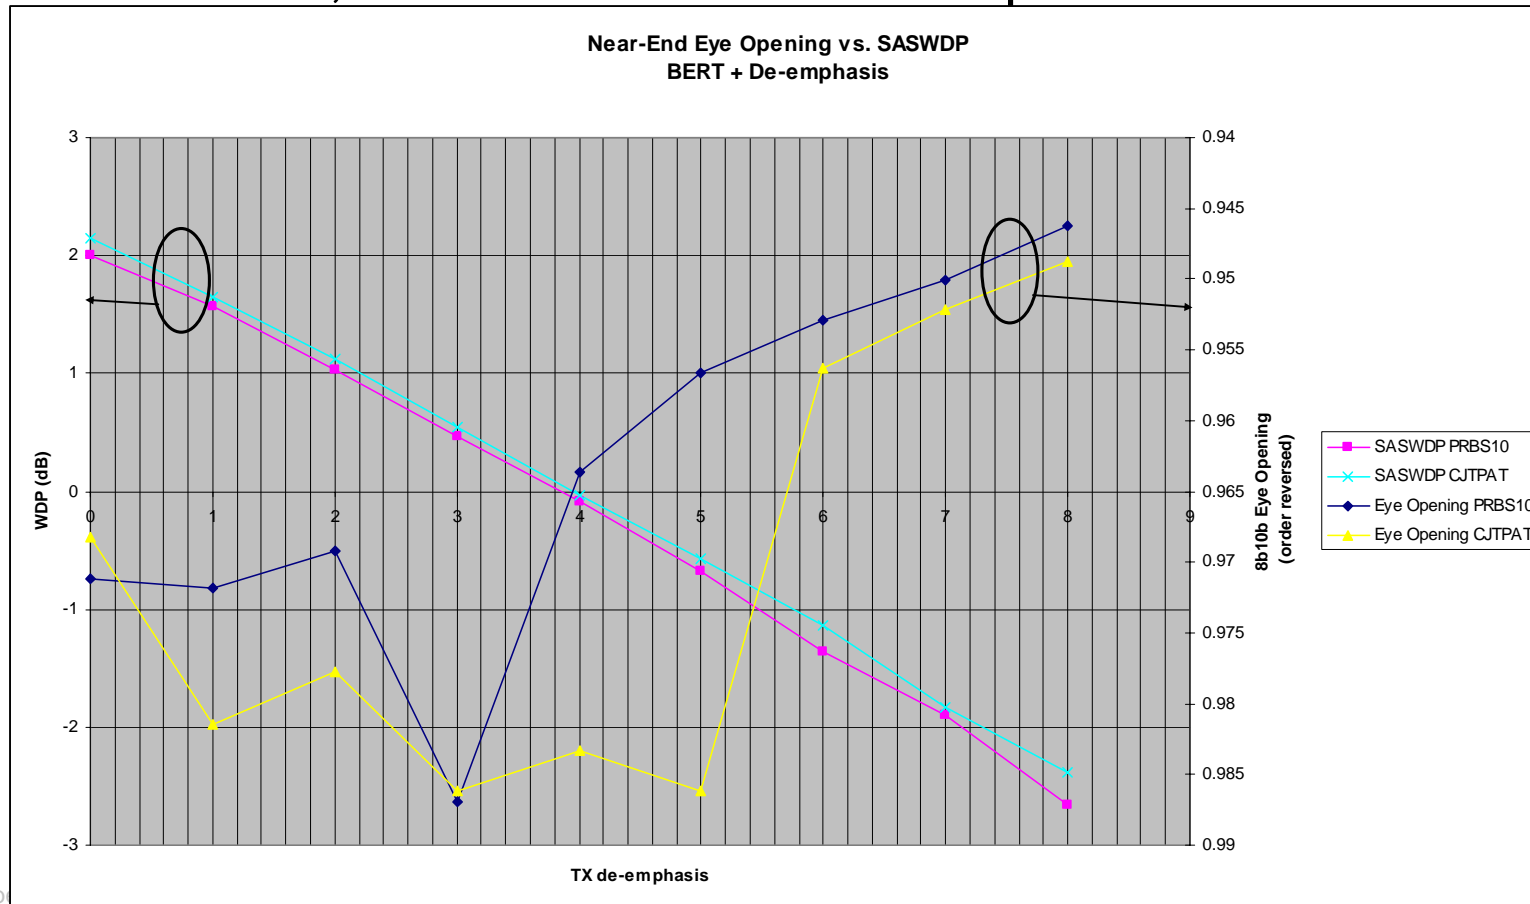
SAS_EYEOPENING vs. SASWDP

- SASWDP vs. SAS_EYEOPENING (PRBS10)



SAS_EYEOPENING vs. SASWDP

- SASWDP vs. SAS_EYEOPENING at **TX NEAR-END**
 - Why is SASWDP WDP dropping? Eye should stay open at near end, or reduce with too much post-cursor.





SAS_EYEOPENING vs. SASWDP

- As DE increases, the main cursor amplitude drops
 - This explains why eye opening slightly degrades: the small energy beyond Tap 3 is scaled by a smaller main cursor

TX_PE=0dB


TX_PE=8dB




SAS_EYEOPENING vs. SASWDP

- xWDP gives strange results:
 - The original script reduces WDP (improves) as PE increases
 - Gain/xMA not constant – effect of CR? The amplitude after scaling by Gain/xMA is larger for PE=8dB, explaining the better SNR

TX_PE=0dB




TX_PE=8dB




SAS_EYEOPENING vs. SASWDP

- If we fix the amplitude extraction (xMA) to use the main cursor's
 - Results are now more constant (only verified for one case)
 - This is what we should expect

TX_PE=0dB



TX_PE=8dB

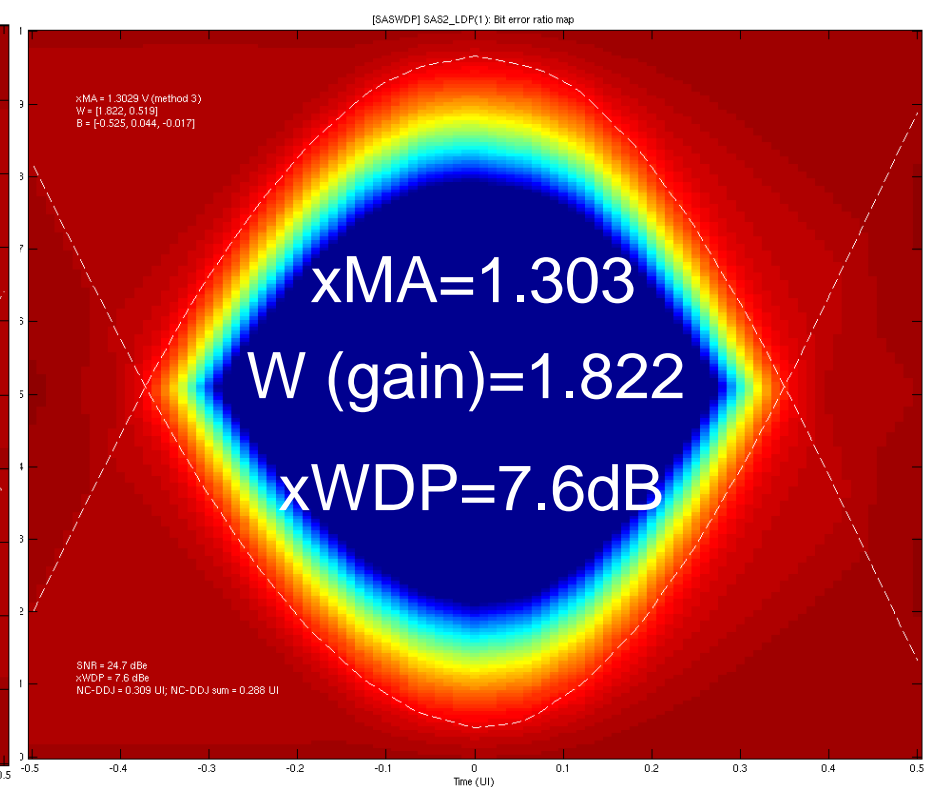



SAS_EYEOPENING vs. SASWDP

- If we fix the amplitude extraction to use the main cursor, but using SASWDP Clock recovery to set it's location
 - Now xWDP increases with increased TX_PE
 - -Due to optimization loop of DFE?

TX_PE=0dB

TX_PE=8dB

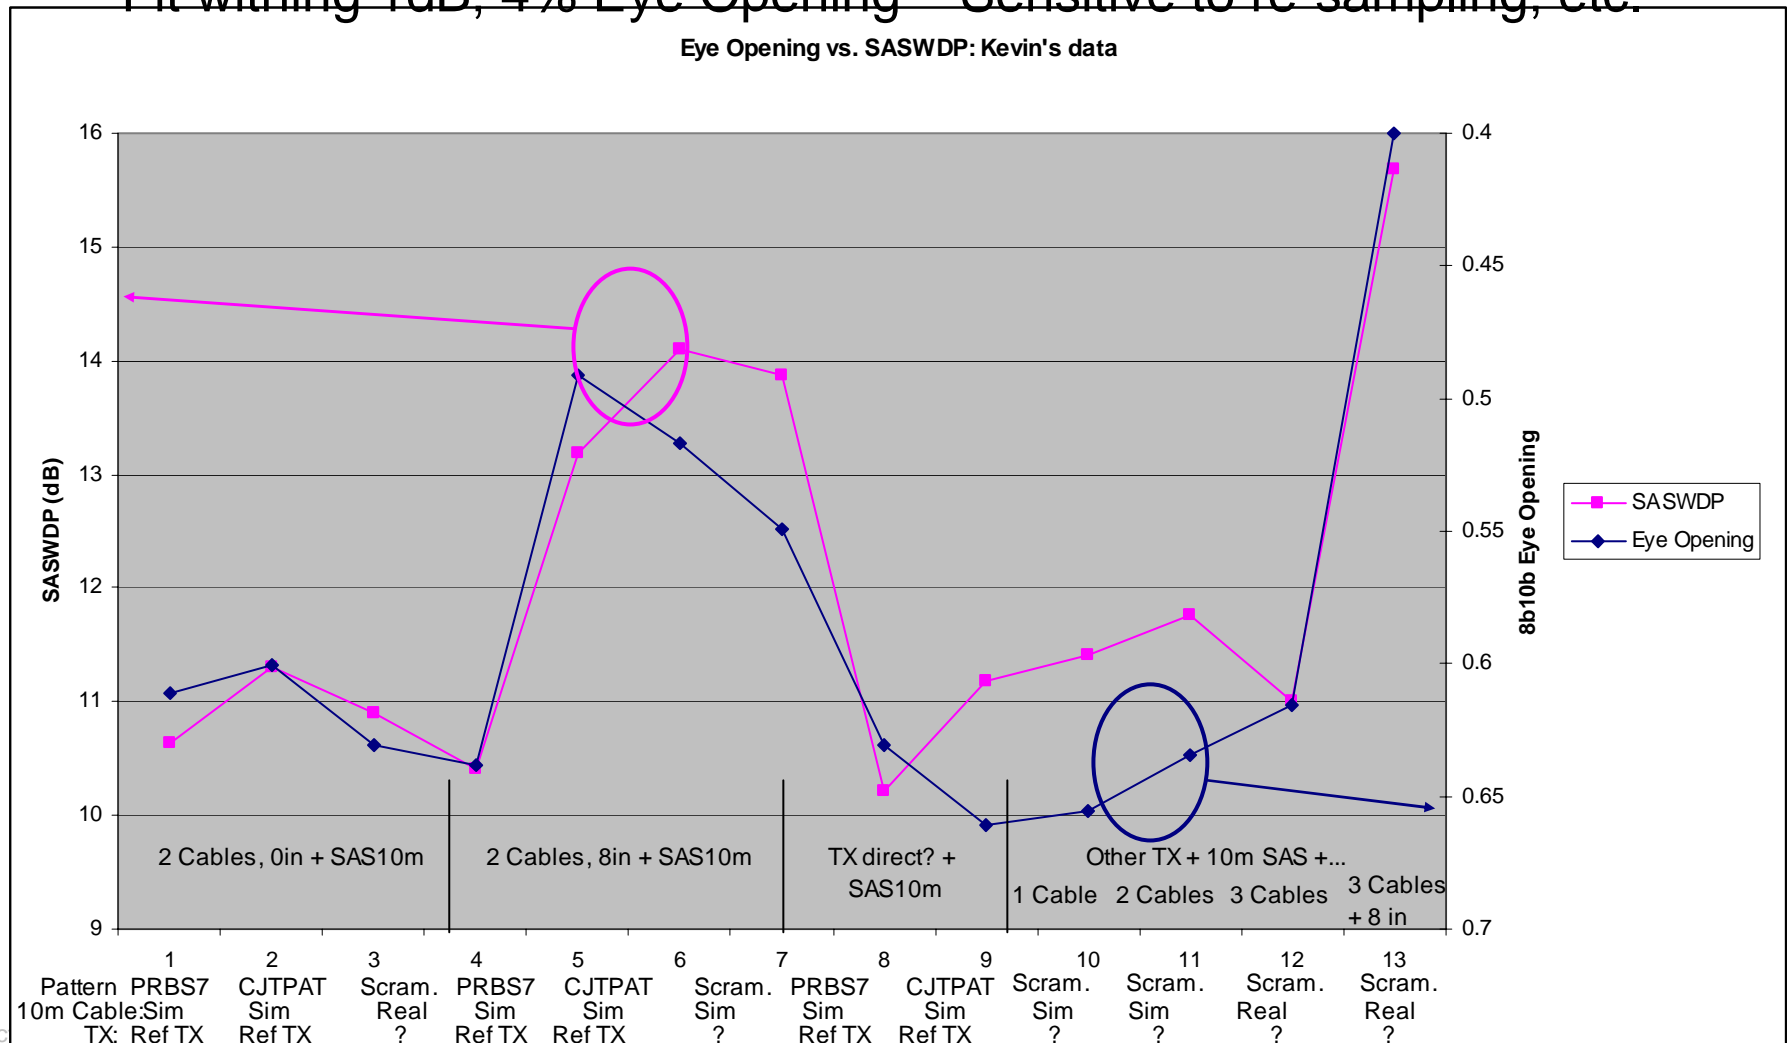


SAS_EYEOPENING vs. SASWDP

- Thus using EYE_OPENING's main tap makes SASWDP results more consistent for processing near-end data
- Near-end data is not what we want to process
 - but that raises questions as whether SASWDP could be useful to process anything but results from a very well-defined environment
 - In theory, we should be able to qualify the TX at the near-end with the same tool

SAS_EYEOPENING vs. SASWDP

- Reprocessed some of Kevin's results
 - Fit withing 1dB, 4% Eye Opening – Sensitive to re-sampling, etc.

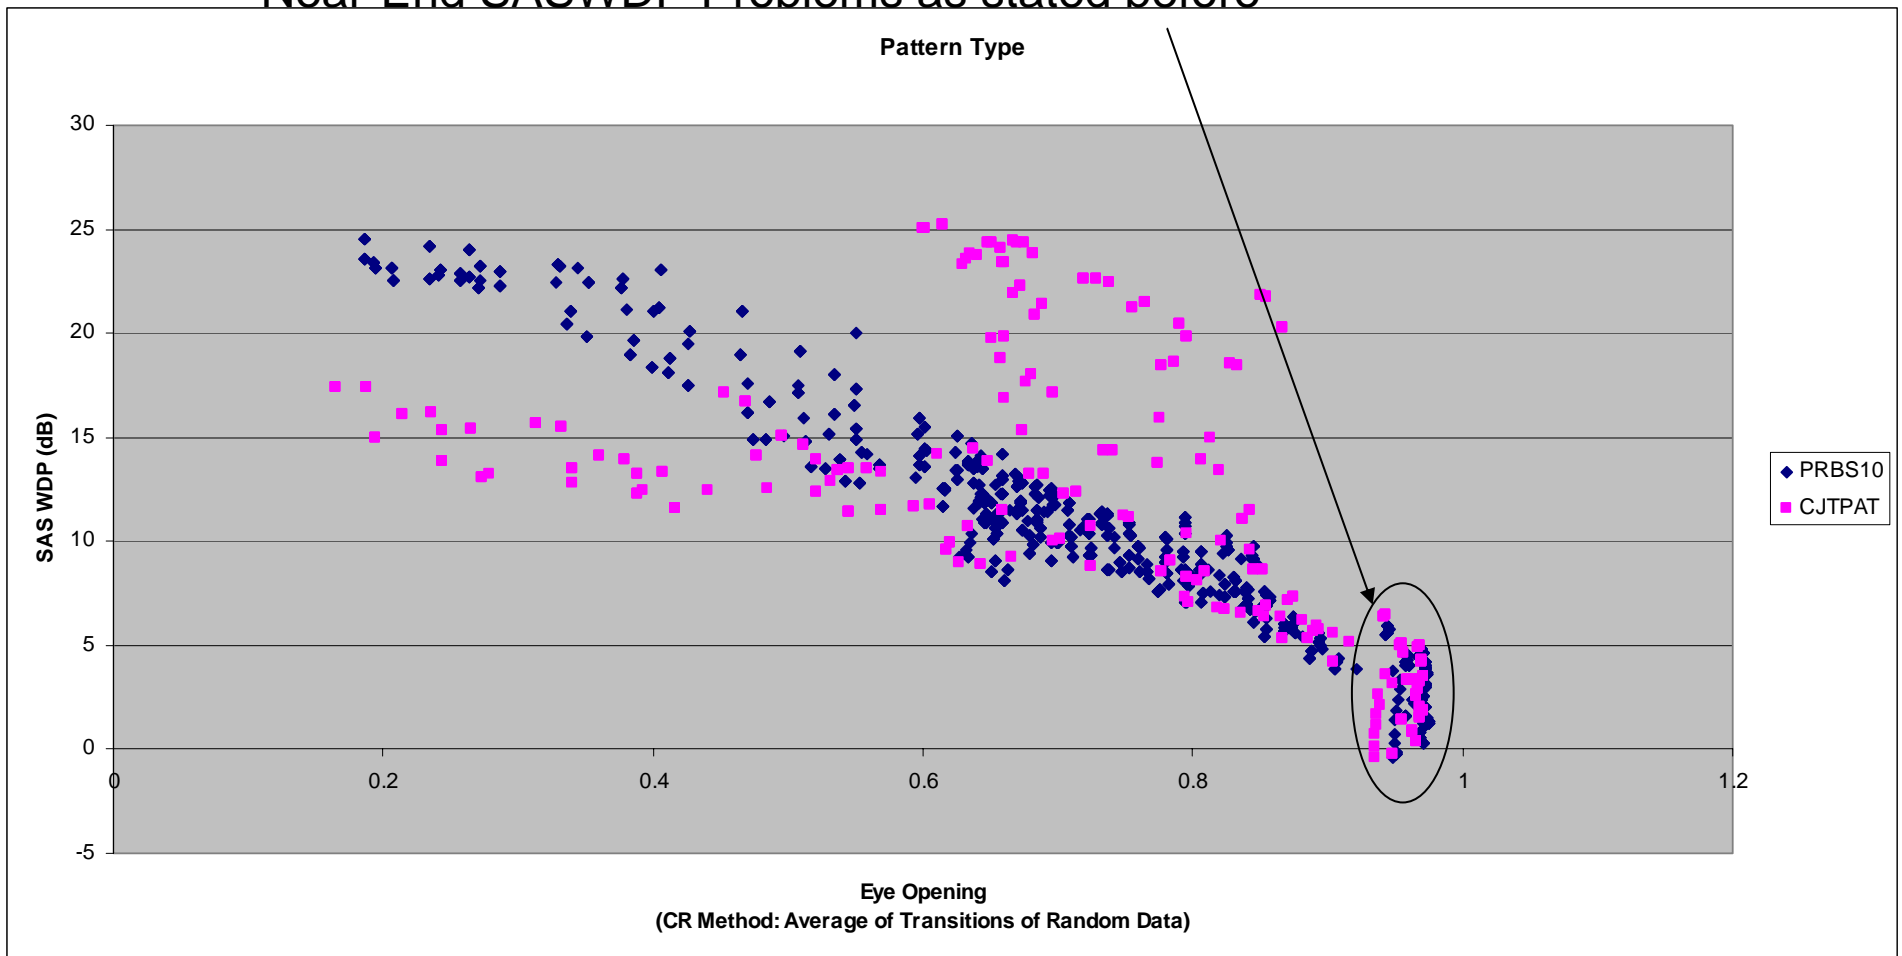


SAS_EYEOPENING vs. SASWDP

- New results generated:
 - 1570 cases (x3 CR Methods for Eye Opening)
 - Synthesized TX (1458 cases x 3 CR Methods)
 - RFT = 41ps, 69ps and 6ps
 - DCD = 40%, 50%, 60%
 - Channels = SAS 10m, 6m, 0.5m, HP24, HP26, PMC Stress
 - Patterns = CJTPAT (Kevin), PRBS10 (PMC), PRBS10 (Kevin)
 - TX De-emphasis = 0, 3, 6 dB
 - TX Zout = 85 Ohms, 100 Ohms, 115 Ohms
 - CR Methods (Eye Opening only) = Peak, 8b10b Transition Average, Random Transition Average
 - Kevin's TX + Waveform (SAS 10m) – 1 case x 3 CR methods
 - Measured BERT Near End (108 cases x 3 CR)
 - Channels = SAS 10m, 6m, 0.5m, HP24, HP26, PMC Stress
 - Patterns = CJTPAT (PMC Lab), PRBS10 (PMC Lab)
 - BERT De-Emphasis (Real) = 0 to 8 dB by 1 dB
 - Measured BERT + 10m SAS Cable + Extra connectors Far End (3 cases x 3 CR)
 - Channel = As is and extra 0.5 m SAS Cable (simulation)
 - Pattern = PRBS10 (PMC Lab)
 - BERT De-Emphasis (Real) = 0 dB and 1dB for Channel as-is

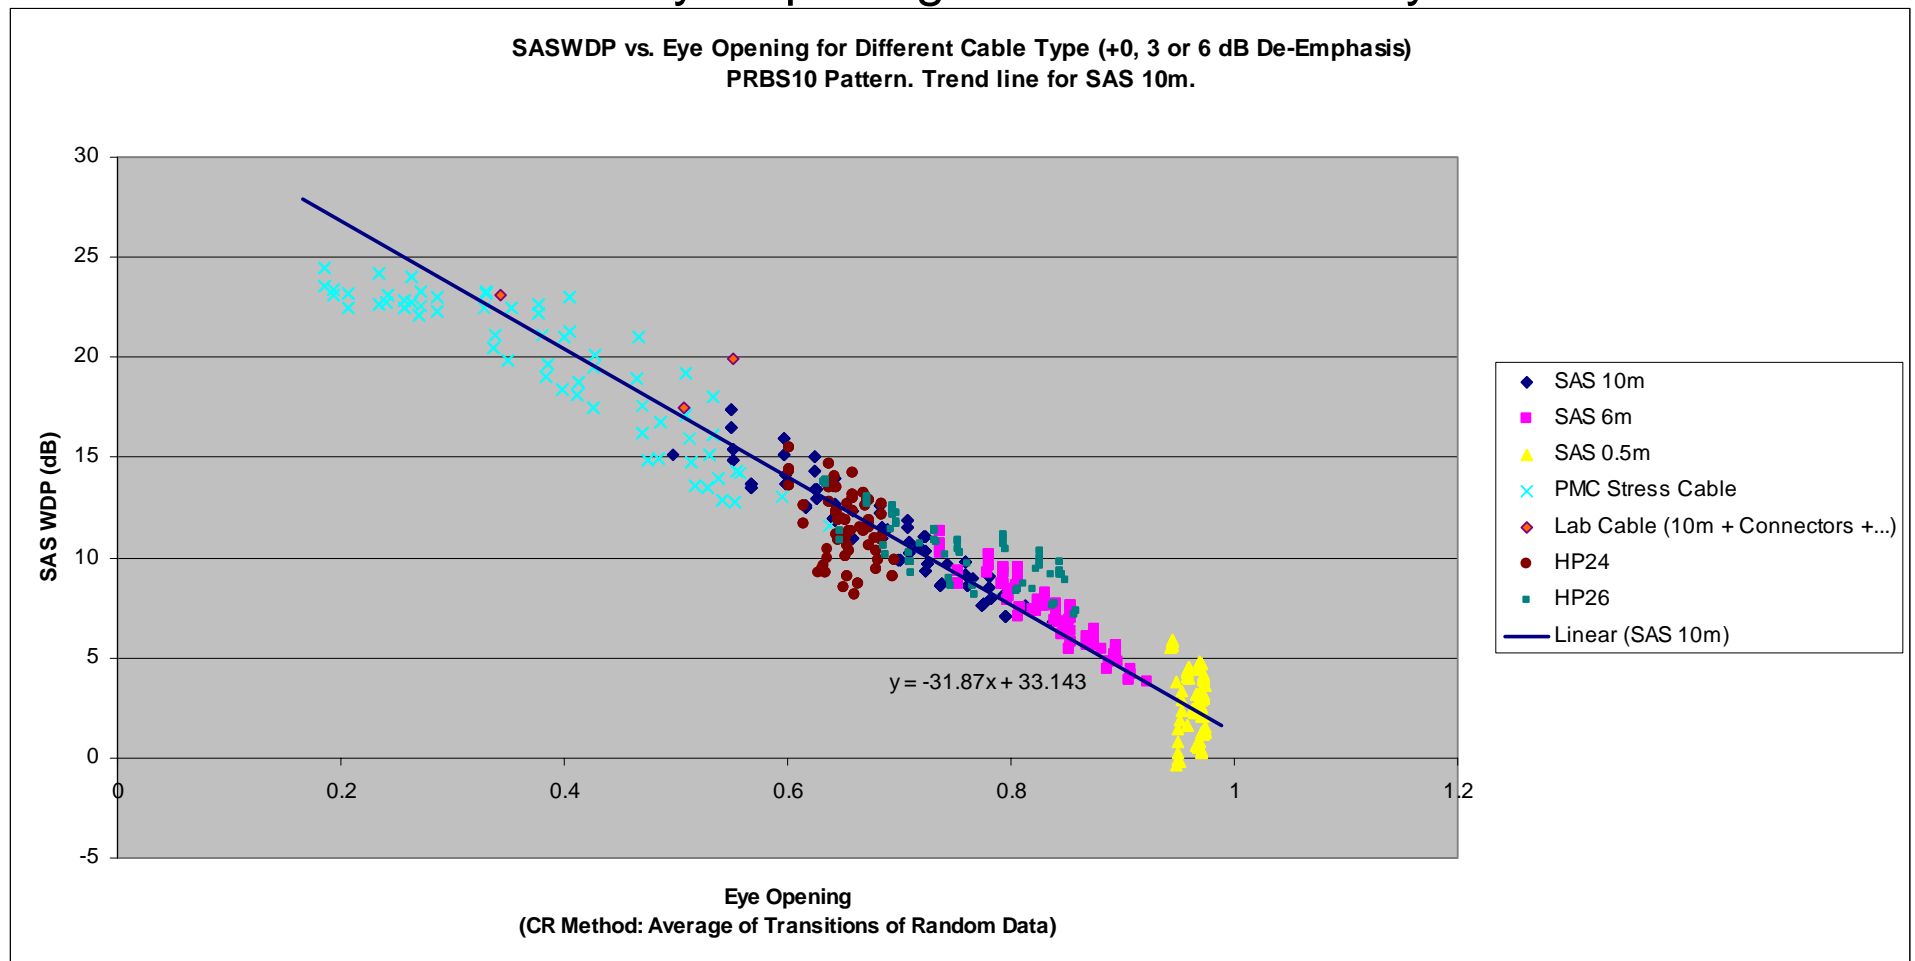
SAS_EYEOPENING vs. SASWDP

- Total Results for Average Transition point from Random Clock Recovery
 - Pattern-dependency of SASWDP very clear
 - Near-End SASWDP Problems as stated before



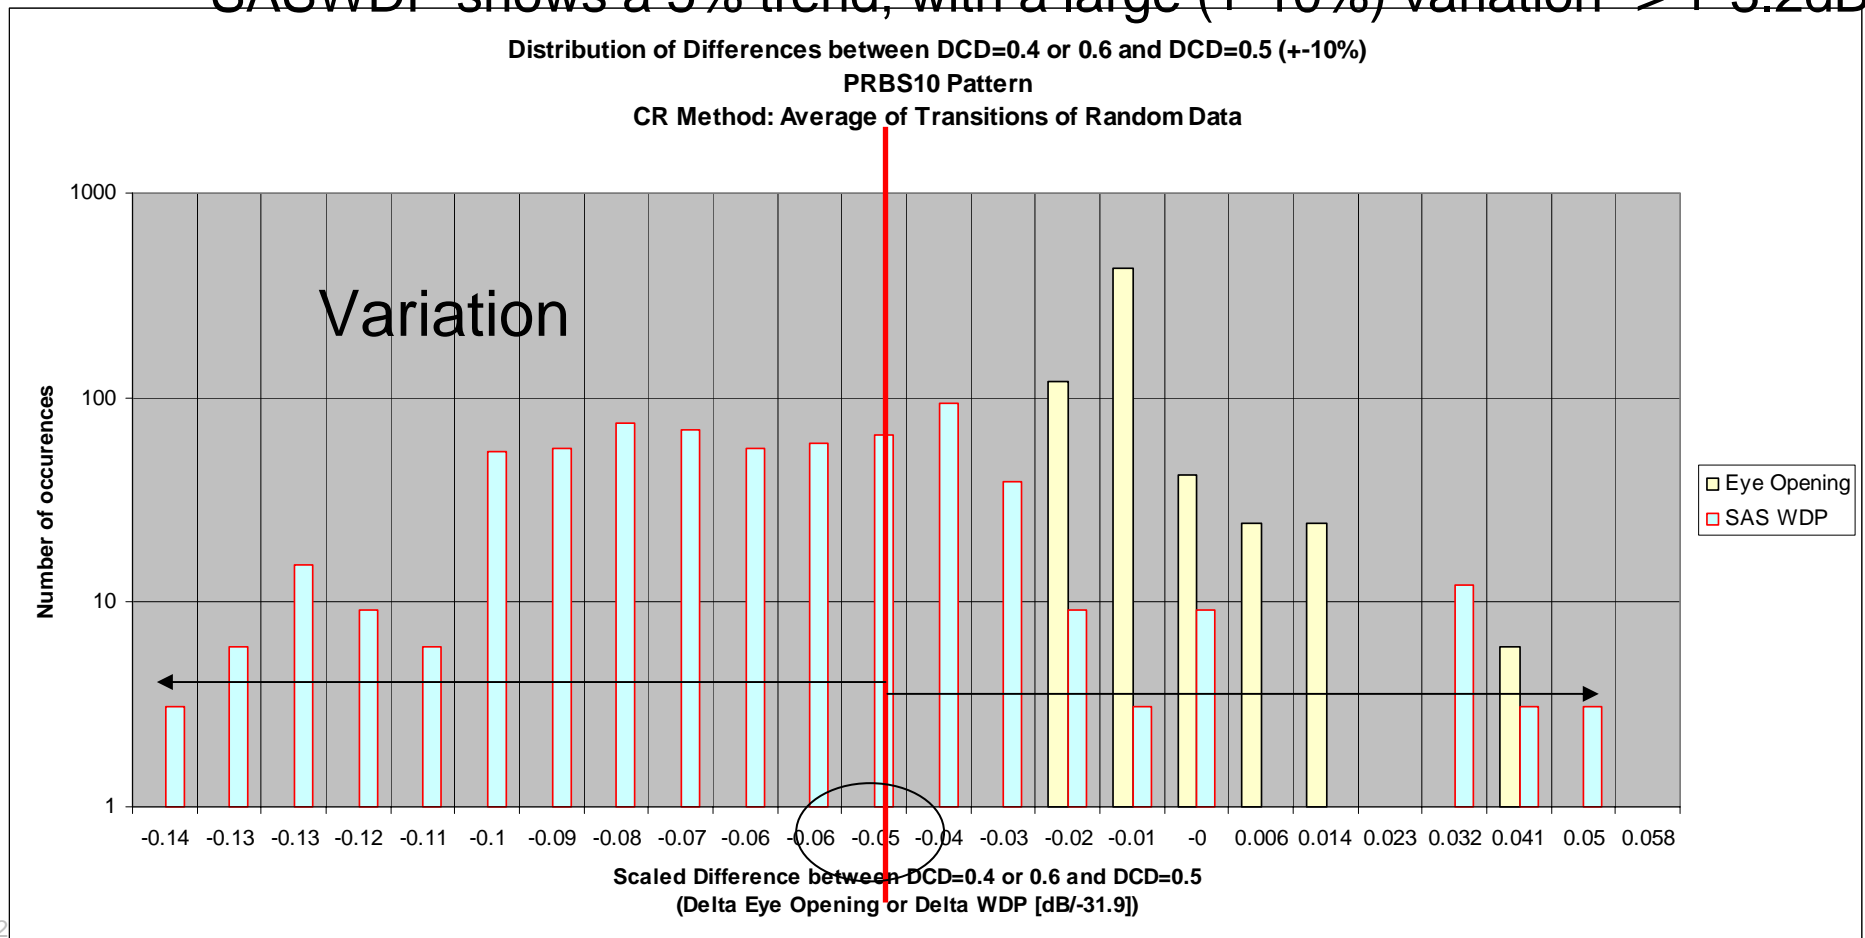
SAS_EYEOPENING vs. SASWDP

- Limiting the Pattern to PRBS10
 - 10m and 6m channels give linear correlation
 - Conversion from Eye Opening to SASWDP is easy



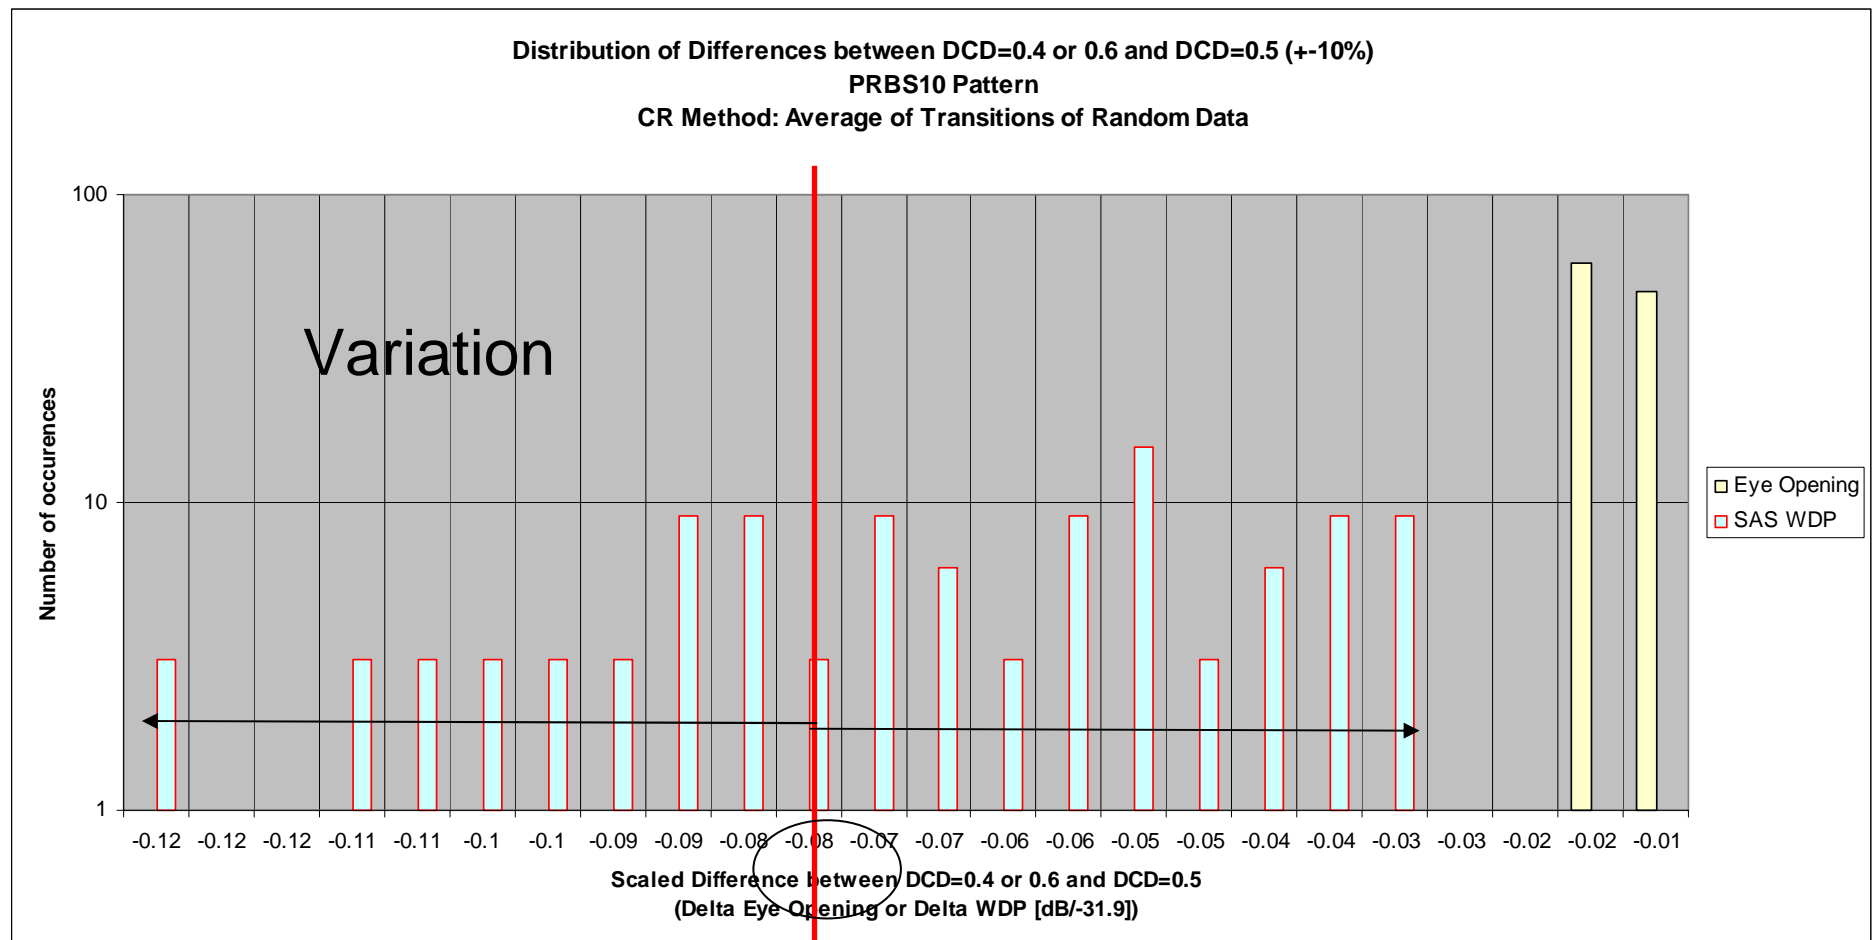
SAS_EYEOPENING vs. SASWDP

- Result Variation Histogram vs. DCD: 0.4 and 0.6 vs. 0.5
 - Eye Opening shows about 0 trend, with a small variation (+-3%)
 - SASWDP shows a 5% trend, with a large (+-10%) variation -> +-3.2dB




SAS_EYEOPENING vs. SASWDP

- Even limiting to SAS 10m, the variation is still large vs. the average for SASWDP



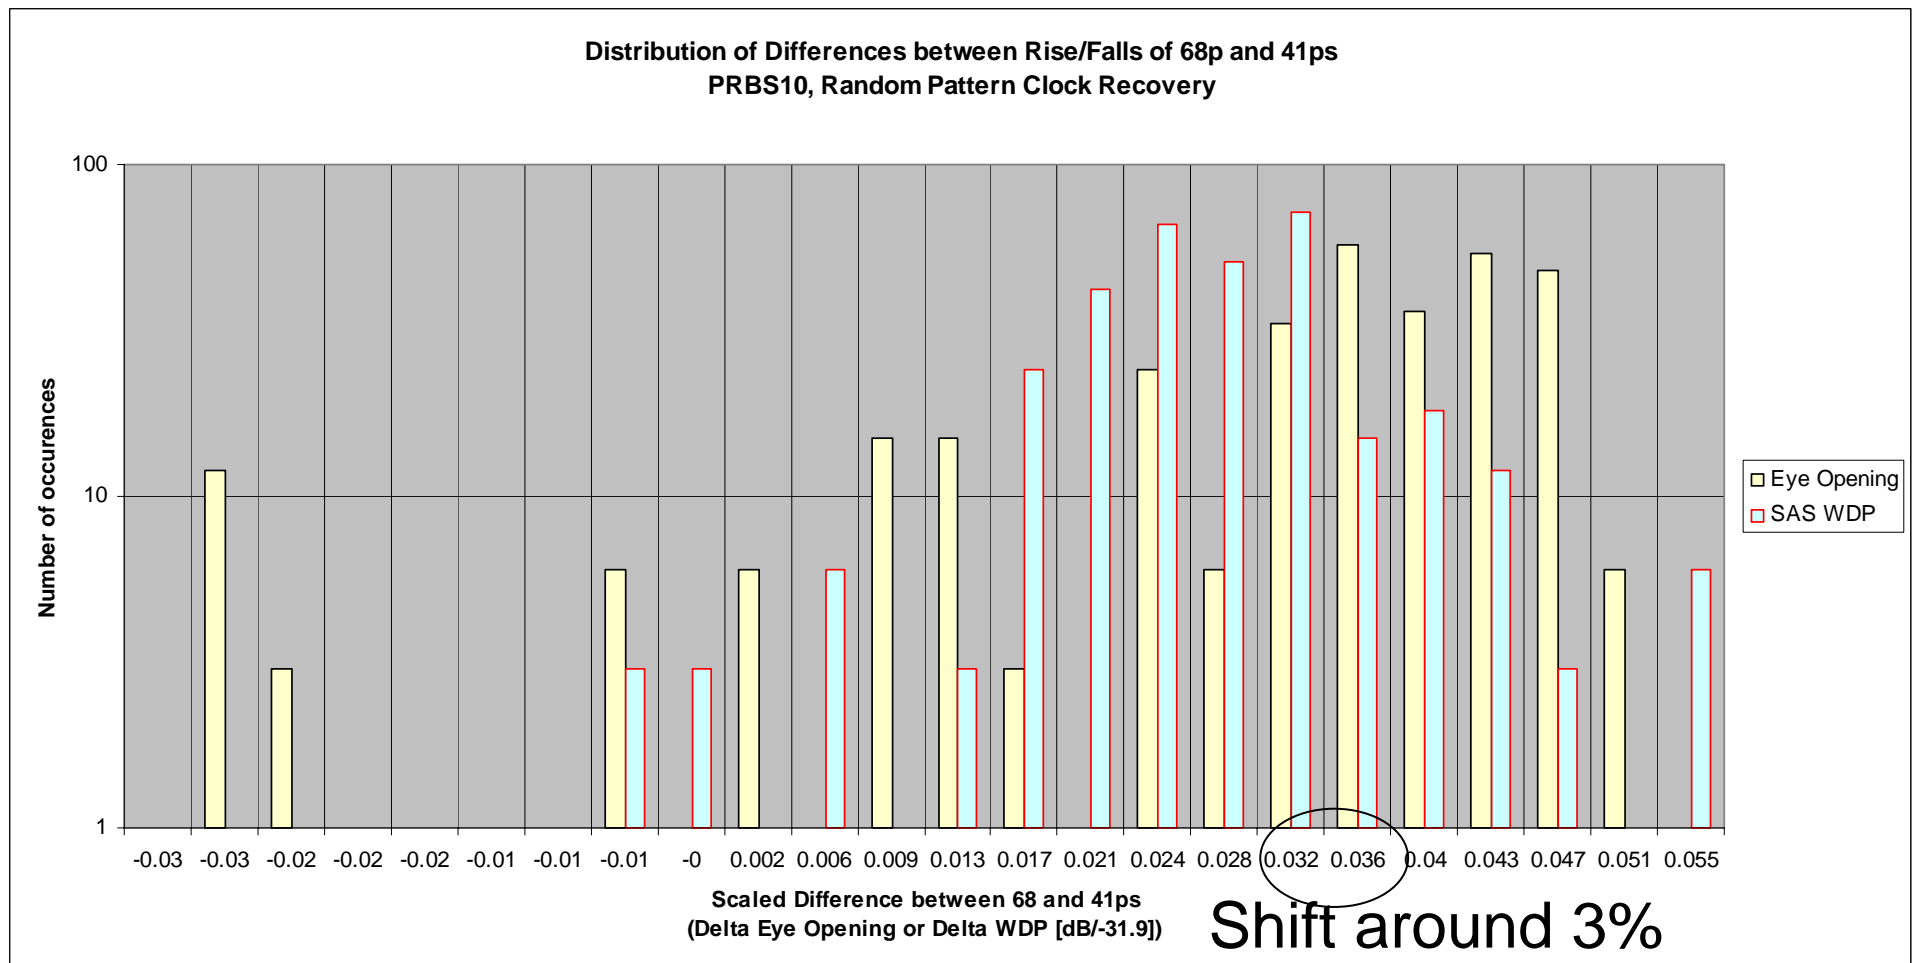
SAS_EYEOPENING vs. SASWDP

- Result Variation Histogram vs. DCD: 0.4 vs .06
 - Eye Opening is equally sensitive to 0.4 and 0.6, SASWDP varies without a clear trend



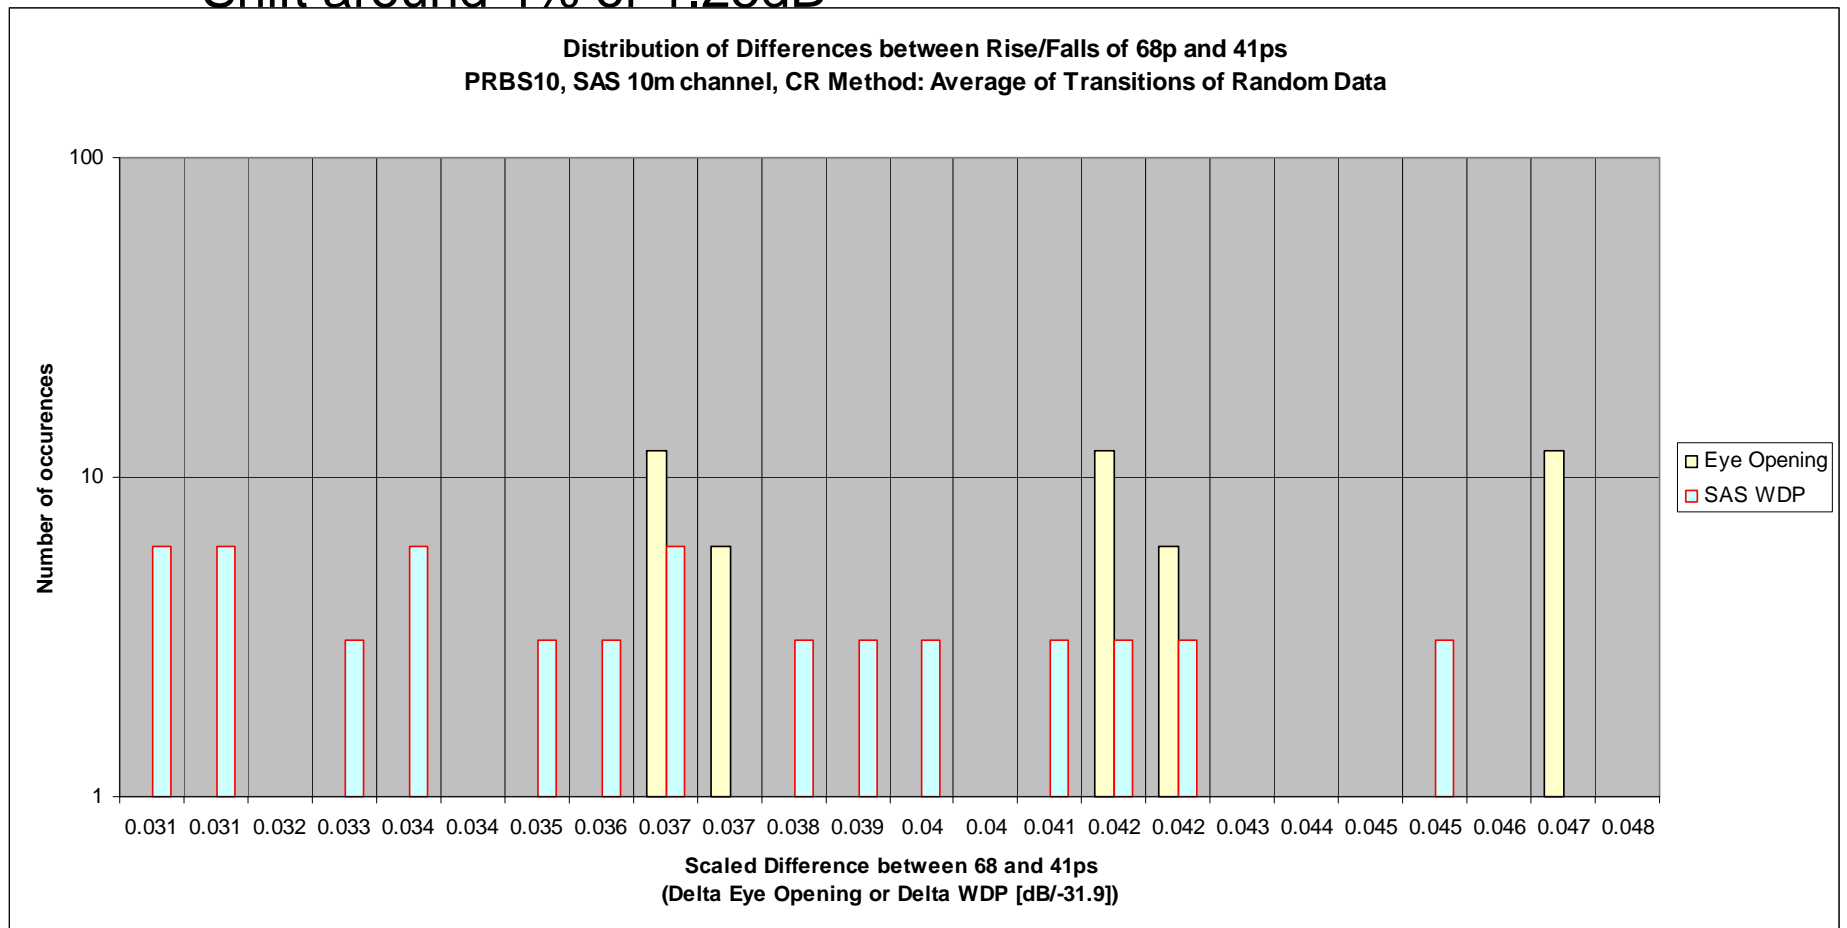
SAS_EYEOPENING vs. SASWDP

- Result Variation Histogram vs. Rise/Fall Times
 - Both are lightly sensitive, which is expected



SAS_EYEOPENING vs. SASWDP

- Limiting to SAS 10m and PRBS10
 - Shift around 4% or 1.25dB



SAS_EYEOPENING vs. SASWDP: TX Impedance



Enabling connectivity. Empowering people.

- Neither are sensitive to TX impedance
 - Absolutely no variation seen in Eye Opening nor SASWDP
 - Maybe this is a simulation issue?
 - However, Main Pulse Amplitude was scaled as expected
 - Return loss may be too small to cause significant reflections in those channels
- The measurement methodology is also pretty insensitive to it (measure TX, then pass through an impulse response)

SASWDP vs. SAS_EYEOPENING: Far-end Measurements Orthogonality



Enabling connectivity. Empowering people.

Specification	Impacts SASWDP result	Impacts SAS_EYEOPENING result
RJ	No (requires manual averaging)	Not Confirmed No (averages itself)
BUJ	No (requires manual averaging) Not Confirmed	No (averages itself)
	DCDJ	Yes, but large variation
	SSCJ	N/A
DDJ	Yes, that's the goal	Yes, that's the goal
Amplitude	Affects W parameter, WDP unchanged	Affects Main tap, Eye Opening Unchanged
De-Emphasis	Yes, that's the goal	Yes, that's the goal
Rise/Fall Times	Yes, that's the goal	Yes, that's the goal
Pattern type	Yes (randomly)	No

SAS_EYEOPENING vs. SASWDP: Summary

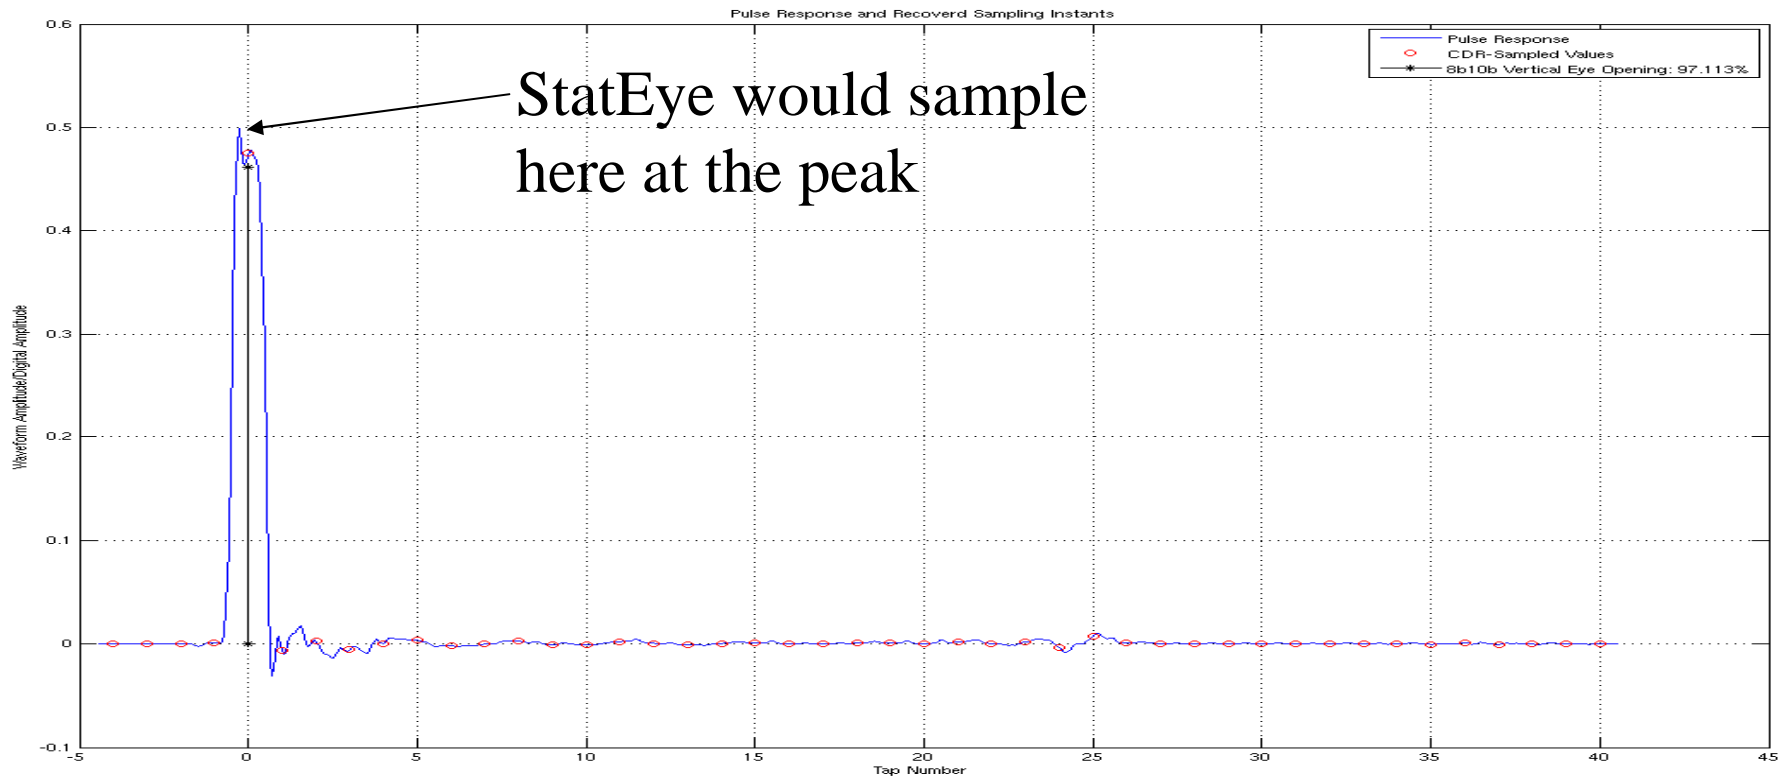
- Good Correlation found between SASWDP and SAS_EYEOPENING
 - Best for SAS 10m Channel, PRBS10 sequence (IDLE is about 2.99dBe below [3.12 for 8b10b CR]– to be confirmed)
 - $xWDP = EyeOpening_{8b10b} * (-31.8696) + 33.14305$ (Random CR)
 - $xWDP = EyeOpening_{8b10b} * (-34.082) + 34.9171$ (8b10b CR)
 - 13 dBe = 64.3%; 15dBe = 58.4%
 - Fits with 1-sigma error of ~0.9 dB
- SASWDP Affected by DCD, Eye Opening much less
 - Delta of ~ +-3dBe for between +-10% DCD
- SASWDP and SAS_EYEOPENING lightly affected by Rise-Fall Times
 - about 4% Eye Opening for min to max. rft spec
 - This is about 1 to 1.5dBe for SASWDP
- Both SASWDP and SAS_EYEOPENING insensitive to TX output impedance
 - Simulation issue? Amplitude scaled however...

StatEye v5.080111 Clock Recovery For Statistical Simulations (stateye.py)



Enabling connectivity. Empowering people.

- The Pulse Response is extracted from the Channel and TX
- The Peak of the Pulse Response is used as the sampling instant
 - This can lead to sensitivity to ripple



StatEye v5.080111 Clock Recovery For Statistical Simulations (stateye.py)



Enabling connectivity. Empowering people.

- Code excerpts:

```
stateye.py, loadStep:  
    self.nomOffset = find(self.pulse==max(self.pulse))[0]
```

Further down, when the actual DFE coefficients are extracted, there are actually comments about this. But all they do, is use this nominal offset to find cursors positions

```
stateye.py, create2TapFIR  
# clearly we need to include here the proper algorithm for finding the optimum sampling point!!  
for cursor in range(noDFEtaps) :  
    self.dfeCoef += [ abs( self.pulse[self.cursor2index(cursor+1, 0)] ) ]  
    print 'Extracting cursor %d, found %0.3f'%(cursor+1, self.dfeCoef[-1])  
  
def cursor2index(self, cursor, sweep) :  
    return( int( round( (cursor+sweep) * self.nomUI + self.nomOffset ) ) ) )
```

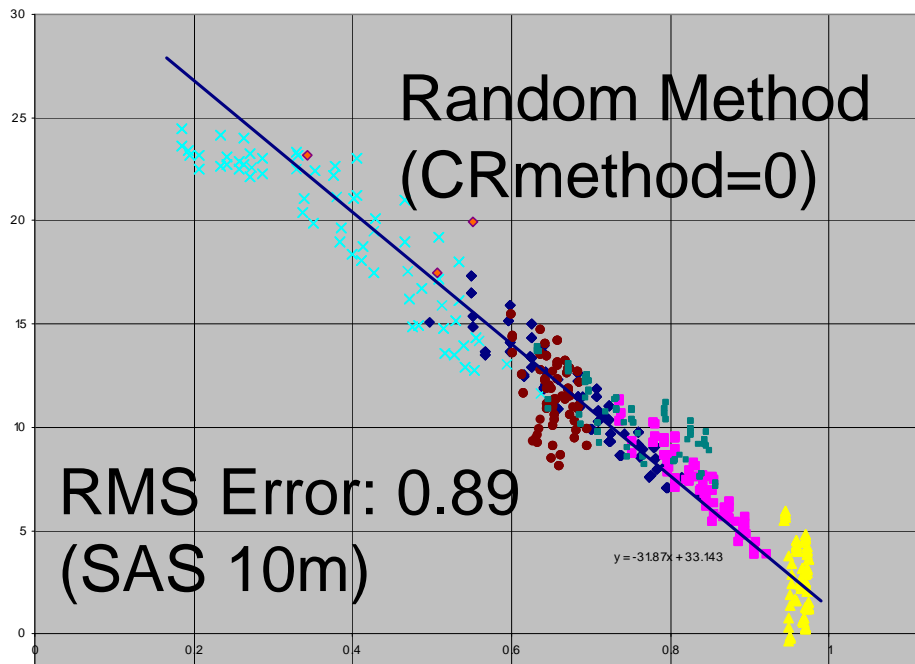
SAS_EYEOPENING Clock Recovery

- Evaluated 3 methods
 - Average Transition Location of Random Pattern (CRmethod=0)
 - The one used so far
 - Peak of Pulse Response (CRmethod=1)
 - The one used in StatEye
 - Average Transition Location of 8b10b Pattern (CRmethod=2)
 - Uses the correlation matrix from an 199800-bits long IDLE pattern

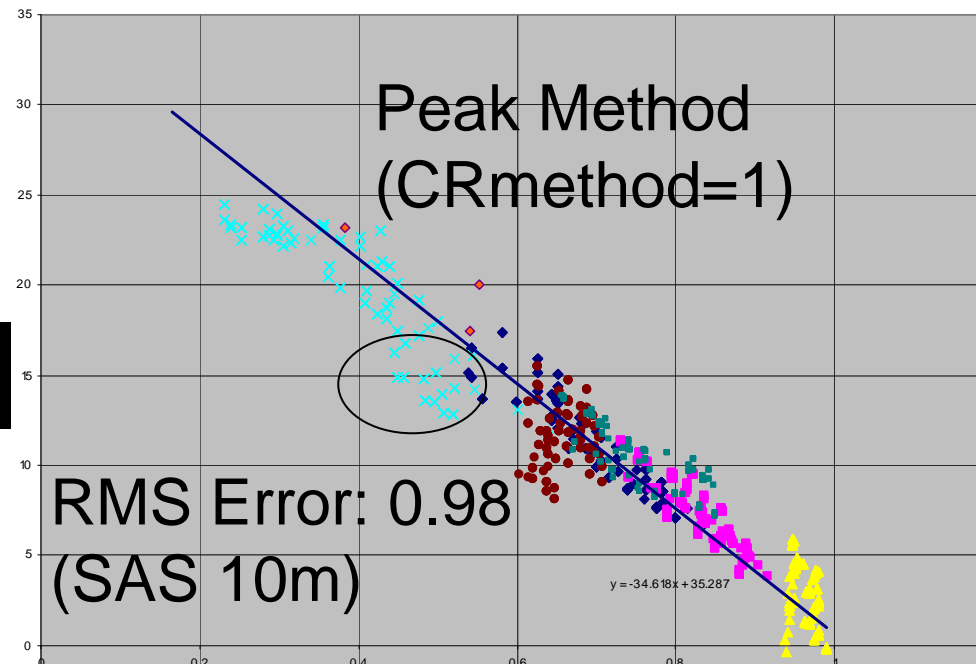
SAS_EYEOPENING Clock Recovery

- All methods give similar results
- Peak Method gives slightly worst performance
 - Likely sensitive to ripple in the pulse response
 - Does not fit as well with SASWDP, but still not bad

SASWDP vs. Eye Opening for Different Cable Type (+0, 3 or 6 dB De-Emphasis)
PRBS10 Pattern. Trend line for SAS 10m.

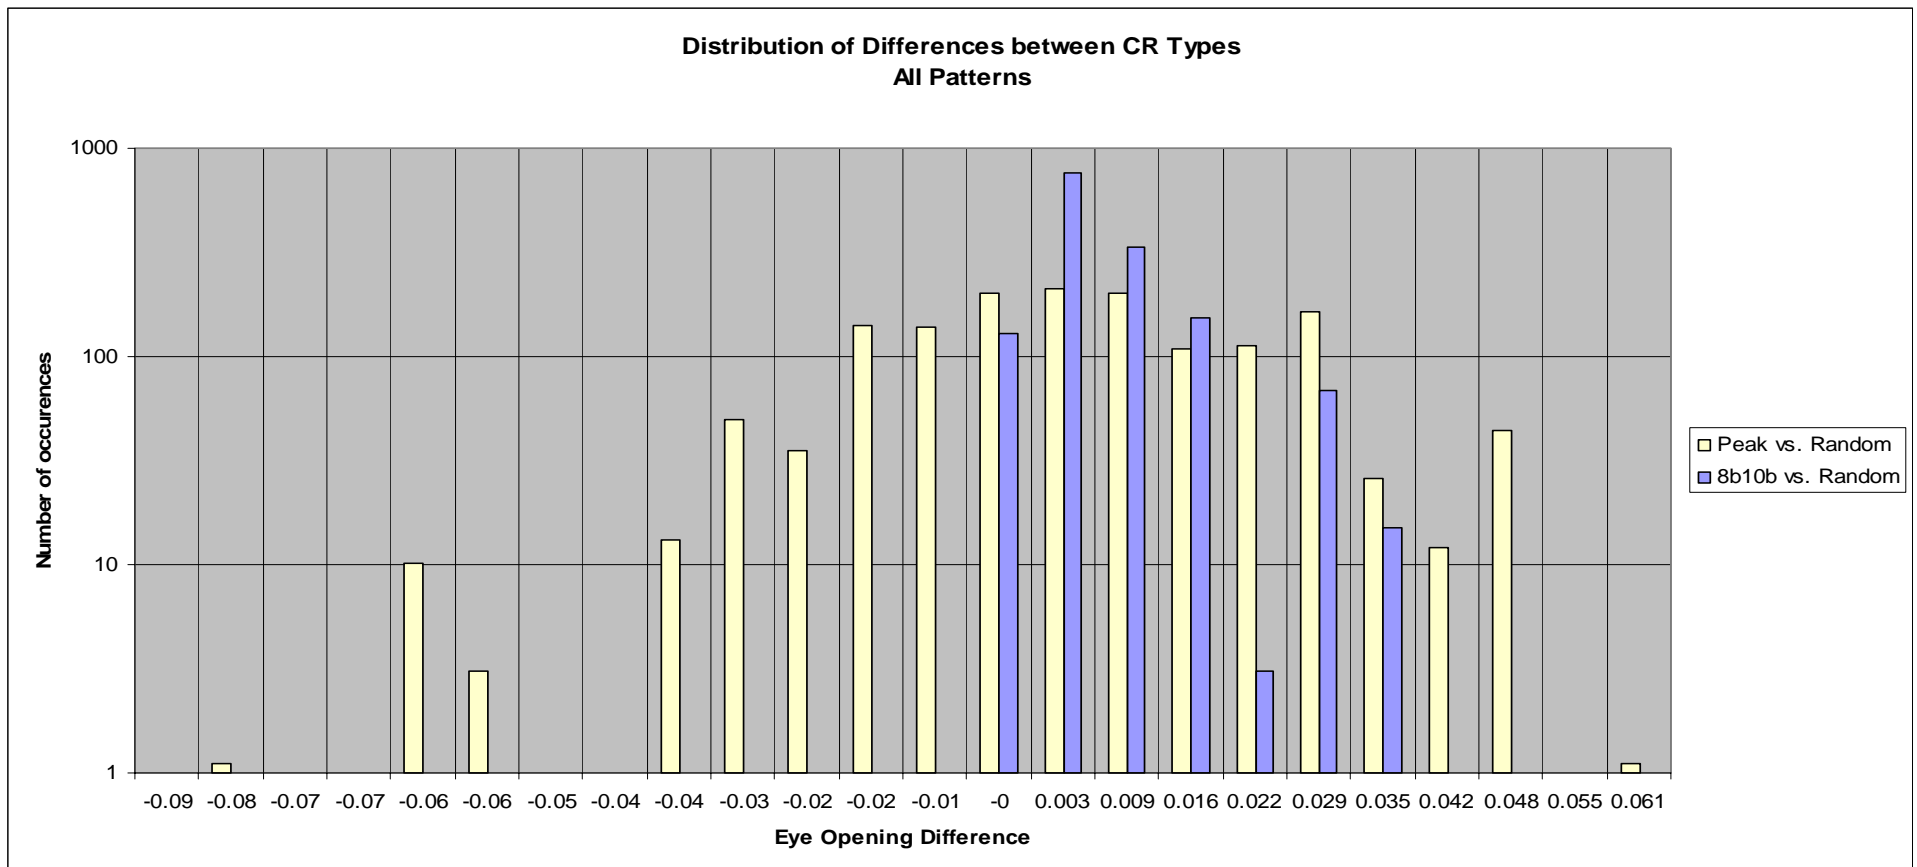


SASWDP vs. Eye Opening for Different Cable Type (+0, 3 or 6 dB De-Emphasis)
PRBS10 Pattern. Trend line for SAS 10m.



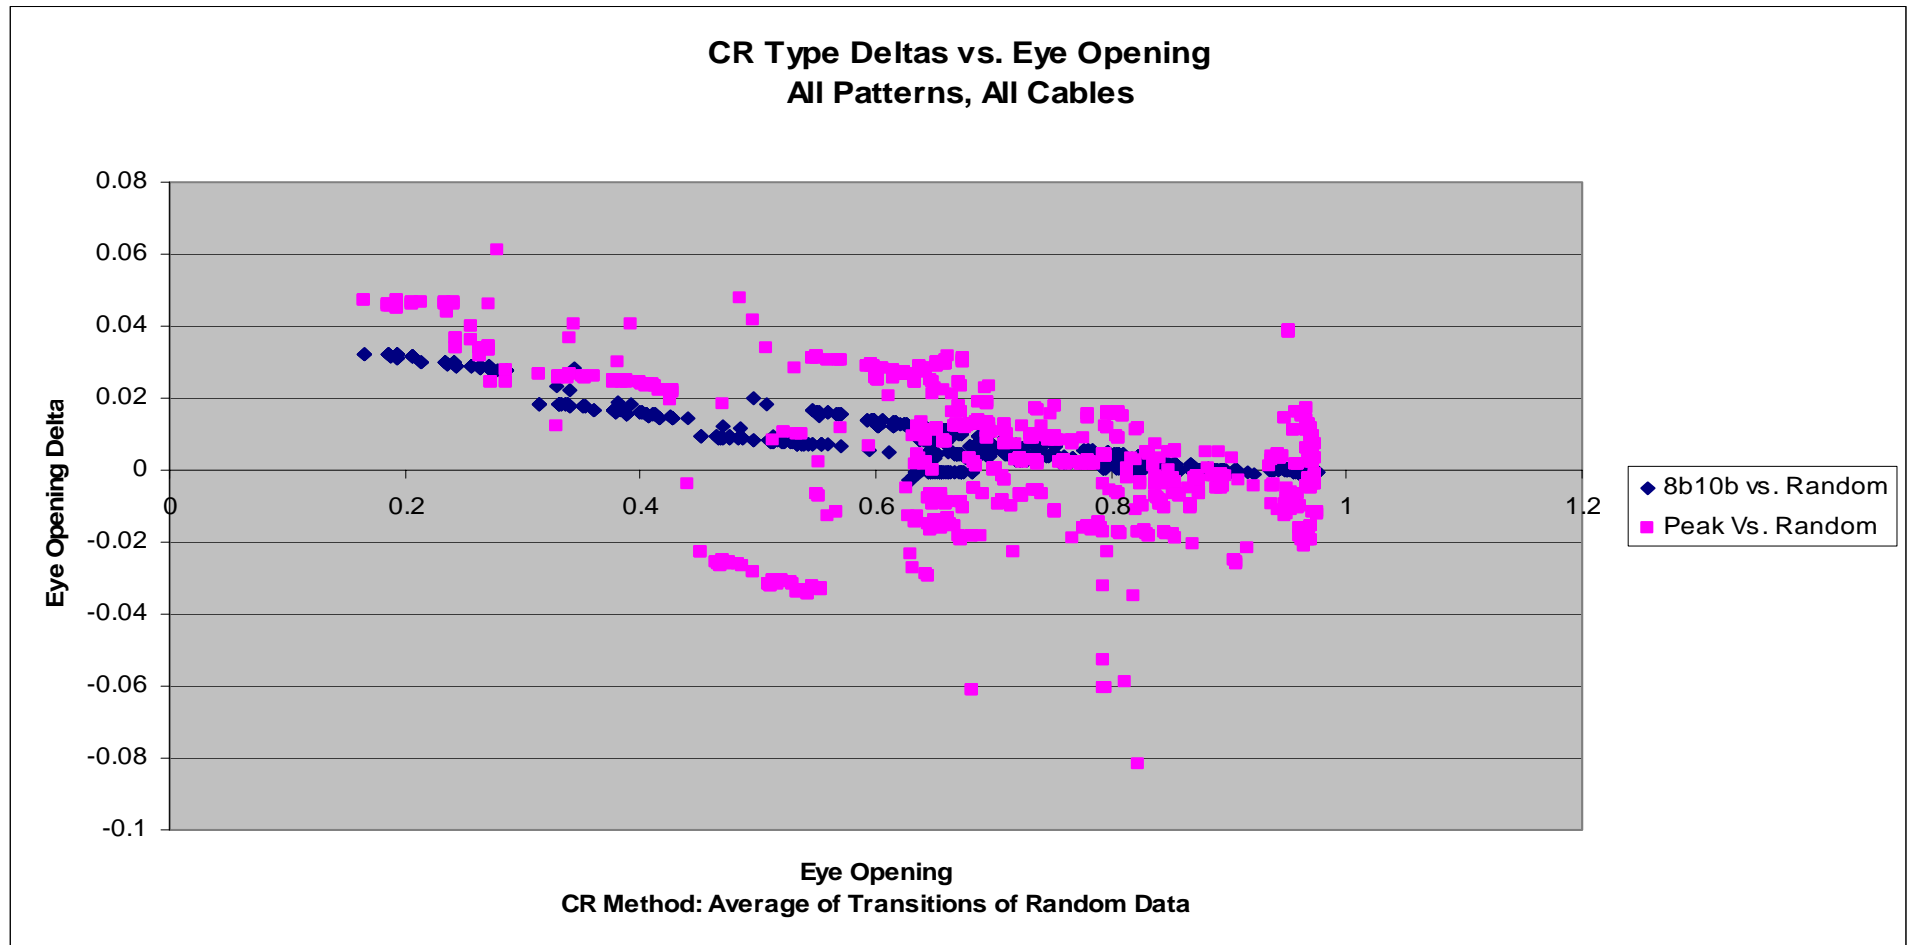
SAS_EYEOPENING Clock Recovery

- 8b10b Method Very Similar to Random
 - 0 to 3.5% delta
- Pulse Peak Method has more Variation
 - -8 to 6%




SAS_EYEOPENING Clock Recovery

- 8b10b vs. Random difference correlated to channel length (i.e. Eye Opening)
- Peak Method varies more (+6/-8%)



SAS_EYEOPENING Clock Recovery

- Even if only PRBS10 & SAS 10m are considered
 - Peak Method Varies $\pm 4\%$



SAS_EYEOPENING Clock Recovery

- Conclusion
 - Using Average of Transitions of Random or 8b10b recommended
 - 8b10b makes more sense vs. expected traffic
 - Results should still match OK with StatEye
 - To be confirmed

- Independent to pattern. CJTPAT can be used.
- Little dependence to rise/fall times
 - This is somewhat part of the DDJ
- Low sensitivity to DCD
- Likely some sensitivity to TX impedance (?)
 - This is also somewhat part of the DDJ
- Does not support SSC
 - Same as SASWDP
- Can be used both at near and far ends
 - Will cover cables from 0 to 10m

SAS_EYEOPENING

Operational Features



Enabling connectivity. Empowering people.

- Can self-correct digital errors
 - Can use data recovered from a closed eye
 - Can thus be used directly from simple time/value scope captures
 - Can re-align data
- Does not require a periodic capture
- Fast & simple to use
 - Long sequences can be used:
 - 120 kBits in ~3 minutes, including correcting > 6k digital errors in 7 passes
 - Insensitive to RJ without pre-averaging

SAS_EYEOPENING.m

Operational Features

- Example run output (from closed-eye recovered data)

```
Iteration 1...
  Number of digital errors: 1038
Iteration 2...
  Number of digital errors: 83
Iteration 3...
  Number of digital errors: 49
Iteration 4...
  Number of digital errors: 25
Iteration 5...
  Number of digital errors: 15
Iteration 6...
  Number of digital errors: 5
Iteration 7...
  Number of digital errors: 2
Iteration 8...
  Number of digital errors: 0
8b10b Eye Opening from fixed CDR = 65.5276% (-3.67151 dB)
Tap Amplitudes:   Main   Tap 1   Tap 2   Tap 3
                  0.225  29.6%  14.1%  11.3%
```

SAS_EYEOPENING Procedure: TX characterization



Enabling connectivity. Empowering people.

- Measure Time/Values at TX near end
- Extract Exact Clock frequency from transitions
 - Or use a synchronous setup
- Extract data from opened-eye
 - Or provide exact sequence if known and align to get the best correlation
- Re-sample values at an integer multiple of the bit rate
 - 16 works fine
- Convolve with Channel
- Send Values and Digital Data to SAS_EYEOPENING
- Compare EyeClosure8b10b with spec.
 - Maybe also DFE 3 taps

SAS_EYEOPENING Procedure: RX characterization



Enabling connectivity. Empowering people.

- Measure Time/Values at RX near end
- Extract Exact Clock frequency from transitions
 - Or use a synchronous setup
- Extract approximative data from closed-eye
 - Or provide exact sequence if known and align to get the best correlation
- Re-sample values at an integer multiple of the bit rate
 - 16 works fine
- Send Values and Digital Data to SAS_EYEOPENING
- Compare EyeClosure8b10b with spec.
 - Maybe also DFE's 3 taps

****Very similar to TX procedure****

- Keep the same methodology as for SASWDP
 - Near-end measurement convolved for TX
 - Far-end measurement for RX
- Change the script from SASWDP to SAS_EYEOPENING
 - Change the spec from WDP & ncDDJ to 8b10b eye opening
 - Add Taps 0 to 3 amplitudes?
- More flexibility (use time/value captures, corrects errors)
- No question about CDR, ncDDJ, SASWDP variability

Suggested Spec. Changes

Modify Section 5.4.7.4.3

- Define the feedback coefficients to have magnitudes less than 0.5 times the peak of the equivalent pulse response
- The taps can be negative (e.g. small channel)
- Can be measured by SAS_EYEOPENING

The reference receiver device includes a 3 tap decision feedback equalizer (DFE) with infinite precision taps and unit interval tap spacing. The reference coefficient adaptation algorithm is the least mean square (LMS) algorithm. The DFE may be modeled at the center of the eye as:

222

Working Draft Serial Attached SCSI - 2 (SAS-2)

26 September 2008

T10/1760-D Revision 14e

$$y_k = x_k - \sum_{i=1}^3 d_i \times \text{sgn}(y_{k-i})$$

where:

y	equalizer differential output voltage
x	equalizer differential input voltage
d	equalizer feedback coefficient
k	sample index in UI


The reference receiver assumes the equalizer feedback coefficients (i.e., d_i) **have** are positive and their magnitudes are less than 0.5. **times the peak of the equivalent pulse response.**

NOTE 28 - For more information on DFE and LMS, see John R. Barry, Edward A. Lee, and David G. Messerschmitt. *Digital Communication - Third Edition*. Kluwer Academic Publishing, 2003. See <http://users.ece.gatech.edu/~barry/digital>.

SAS_EYEOPENING.m can be used to compute the relative amplitude.

Suggested Spec. Changes

A 0.5 m Channel with 3 dB of TX de-emphasis



Suggested Spec. Changes

Modify Table 60

- See 08-330r4 p. 7
 - How can we average a $2^{16}-1$ bit pattern after 8b10b encoding?
 - Is this easy to do with standard test equipment?
- Add a note that significant amount of DCD may impact measurement with this script. Suggest alternate script?
- Can we extract RJ out of SSC? (note g)

Table 60 — Transmitter device signal output characteristics for trained 1.5 Gbps, 3 Gbps, and 6 Gbps at IT and CT

Signal characteristic	Units	Minimum	Nominal	Maximum
Peak to peak voltage if SATA is not supported ^a	mV(P-P)	850		1 200
Transmitter device off voltage ^b	mV(P-P)			50
Withstanding voltage (non-operational)	mV(P-P)	2 000		
Rise/fall time ^c	UI	0.25 ^d		
Reference differential impedance ^e	ohm		100	
Reference common mode impedance ^e	ohm		25	
Common mode voltage limit (rms) ^f	mV			30
Random jitter (RJ) ^{g, j}	UI			0.15 ^h

Maximum eye opening (e.g., 20-74 ps) (Figure 121)

- ^a See 5.4.6.4.5 for measurement method.
- ^b The transmitter device off voltage is the maximum A.C. voltage measured at compliance points IT and CT when the transmitter is unpowered or transmitting D.C. idle (e.g., during idle time of an OOB signal).
- ^c Rise/fall times are measured from 20 % to 80 % of the transition with a repeating 01b or 10b pattern (e.g., D10.2 or D21.5)(see table 237 in 10.2.9.2) on the physical link.
- ^d 0.25 UI is 41.6 ps at 6 Gbps.
- ^e For transmitter device S-parameters characteristics, see 5.4.6.4.2.
- ^f This is a broadband limit. For additional limits on spectral content, see figure 131 and table 61.
- ^g RJ is 14 times the RJ 1 sigma value, based on a BER of 10^{-12} . This test shall be performed with a repeating 01b or 10b pattern (e.g., D10.2 or D21.5)(see table 237 in 10.2.9.2) on the physical link. If the transmitter device supports SSC, then this measurement shall be performed with both SSC enabled and SSC disabled. For simulations based on a BER of 10^{-15} , the RJ specified is 17 times the RJ 1 sigma value.
- ^h 0.15 UI is 25 ps at 6 Gbps.
- ⁱ 0.30 UI is 50 ps at 6 Gbps.
- ^j See 5.4.5.2 for JMD requirements.
- ^k This value is obtained by simulation. It represents the resulting signal output within the reference receiver device (see 5.4.7.4.3) after equalization, when the transmitter device output signal of CJTPAT is transmitted through the reference transmitter test load (see 5.4.2.5). The specific simulation program used (e.g., StatEye from <http://www.stateye.org>) is not specified by this standard.

Suggested Spec. Changes

Modify Table 60

- Add Near-end Measurement spec. for TJ max = 0.25 UI
 - Same as reference transmitter
 - BUJ=0.1 + RJ=0.15
 - Measure with D24.3 Pattern (to remove effect of any TX de-emphasis, reduce ISI impact)
 - Simply TJ
 - JTF transition density nominal
 - or with CJTPAT
 - TJ minus DDJ

Reference TX BUJ=0.1 ←

Table 60 — Transmitter device signal output characteristics for trained 1.5 Gbps, 3 Gbps, and 6 Gbps at IT and CT

Signal characteristic	Units	Minimum	Nominal	Maximum
Peak to peak voltage if SATA is not supported ^a	mV(P-P)	850		1 200
Transmitter device off voltage ^b	mV(P-P)			50
Withstanding voltage (non-operational)	mV(P-P)	2 000		
Rise/fall time ^c	UI	0.25 ^d		
Reference differential impedance ^e	ohm		100	
Reference common mode impedance ^e	ohm		25	
Common mode voltage limit (rms) ^f	mV			30
Random jitter (RJ) ^{g, j}	UI			0.15 ^h

Maximum eye opening (V_{pe}) at IT and CT

^a See 5.4.6.4.5 for measurement method.
^b The transmitter device off voltage is the maximum A.C. voltage measured at compliance points IT and CT when the transmitter is unpowered or transmitting D.C. idle (e.g., during idle time of an OOB signal).
^c Rise/fall times are measured from 20 % to 80 % of the transition with a repeating 01b or 10b pattern (e.g., D10.2 or D21.5)(see table 237 in 10.2.9.2) on the physical link.
^d 0.25 UI is 41.6 ps at 6 Gbps.
^e For transmitter device S-parameters characteristics, see 5.4.6.4.2.
^f This is a broadband limit. For additional limits on spectral content, see figure 131 and table 61.
^g RJ is 14 times the RJ 1 sigma value, based on a BER of 10⁻¹². This test shall be performed with a repeating 01b or 10b pattern (e.g., D10.2 or D21.5) see table 237 in 10.2.9.2) on the physical link. If the transmitter device supports SSC, then this measurement shall be performed with both SSC enabled and SSC disabled. For simulations based on a BER of 10⁻¹⁵, the RJ specified is 17 times the RJ 1 sigma value.
^h 0.15 UI is 25 ps at 6 Gbps.
ⁱ 0.30 UI is 50 ps at 6 Gbps.
^j See 5.4.5.2 for JMD requirements.
^k This value is obtained by simulation. It represents the resulting signal output within the reference receiver device (see 5.4.7.4.3) after equalization, when the transmitter device output signal of CJTPAT is transmitted through the reference transmitter test load (see 5.4.2.5). The specific simulation program used (e.g., StatEye from <http://www.stateye.org>) is not specified by this standard.

Table 64 — Reference transmitter device characteristics at IT and CT

Characteristic	Units	Value
Peak to peak voltage (V _{P-P}) ^a	mV(P-P)	850
Transmitter equalization ^a	dB	2
Maximum rise/fall time ^b	UI	0.41 ^c
RJ	UI	0.15 ^d
BUJ	UI	0.10 ^e

^a See 5.4.6.4.5 for measurement method.
^b Rise/fall times are measured from 20 % to 80 % of the transition with a repeating 01b or 10b pattern (e.g., D10.2 or D21.5)(see table 237 in 10.2.9.2).
^c 0.41 UI is 68.3 ps at 6 Gbps.
^d 0.15 UI is 25 ps at 6 Gbps.
^e 0.10 UI is 16.6 ps at 6 Gbps.

Suggested Spec. Changes

Modify Table 71

- See 08-330r4 p. 13
 - Add note to Measure without the BUJ and RJ
 - Assumes BERT RJ is very small
- Table 223?
- Make stressed receiver a subset of the jitter tolerance test?
 - Use a second mask when SSC is not used?

Table 71 — Stressed receiver device sensitivity test characteristics

Characteristic	Units	Minimum	Typical	Maximum	Reference
Tx data pattern			CJTPAT		Annex A
Tx peak to peak voltage	mV(P-P)		850	800	5.4.6.4.1
Transmitter equalization	dB		2		5.4.6.4.5
Tx RJ	UI	0.15 ^b			5.4.6.4.1
Link dispersion penalty ^{d, e}	dB	13			5.4.7.4.4.8
D24.3 delivered eye opening (Z1) ^e	mV	75		95	5.4.5.4
NEXT offset frequency ^e	ppm	20			
Total crosstalk amplitude ^{e, f}	mV _{rms}	4			

^a 0.24 UI is 41.6 ps at 6 Gbps.
^b 0.15 UI is 25 ps at 6 Gbps.
^c 0.000 22 UI is 0.036 ps at 6 Gbps.
^d Link dispersion penalty is the WDP of the delivered signal computed with Palloc = 15.4 dB.
^e This specification pertains to the delivered signal at IR or CR during the receiver device compliance test. All adjacent phys in the receiver device shall be active with representative traffic with their maximum amplitude and maximum frequency of operation. Additional pseudo-random crosstalk shall be added, if needed, to meet the total crosstalk amplitude specification.
^f Observed with a histogram of at least 1 000 hits.

- Test orthogonality to more effects

To Do

BUJ, DCD Done

- TX impedance Done ?

- Test with more TX and channels

- Try to make output WDP-like

- Re-do for IDLE pattern

- Look at StatEye's CDR Done

- Compare StatEye to Eye Opening for Synthetic TX cases

- Define a min. eye opening spec. for the TX Done

- Define a min & max. eye opening spec. for the RX Done

- Evaluate if a spec. on DFE coefficients is required

- Clarify spec. Done

- Evaluate near-end results convolved with channel Done

- Look at applicability to 1.5 and 3G

To Do

Other considerations



Enabling connectivity. Empowering people.

StatEye v5.2 Clock Recovery For Pattern Recovery (extractsignal.py)




Enabling connectivity. Empowering people.

- Data is filtered by a digital filter
 - Parameters:
 - Alpha (high-pass)
 - Beta (low-pass)
 - Time Step (overall frequency scaling)
- Serves as noise filter
- Can serve as equalization
- Examples (testall.bat) included Alpha=5 and Default Alpha=60
 - Parameter z in testcase.py
 - Beta always 0.1
- Some examples forced Time Step=12.5ps
 - Parameter x in testcase.py
- Further NOTE from Rob Elliott:
 - This is for the extraction from the waveform (extractsignal.py)
 - The part of the code we are interested in for correlation with previous result does not contain this routine

StatEye v5.2 Clock Recovery For Pattern Recovery (extractsignal.py)

- Example Responses:

- $\alpha=5$ or 60
- $\beta=0.1$ (default)
- $t=12.5\text{ps}$
or 25ps



StatEye v5.2 Clock Recovery For Pattern Recovery (extractsignal.py)



Enabling connectivity. Empowering people.

- CR then works on the edges of the « emphasized » input
 - Emphasis based on user-defined parameters & time-step
 - It seems to then filter the transitions further
 - Parameters m and k in cdr.py
 - $\text{period} += [\text{period}[-1] + \text{phaseError}[-1] * k]$
 - $\text{phase} += [\text{phase}[-1] + \text{phaseError}[-1] * m + \text{nperiod}[-1] * \text{period}[-1]]$
- An Iterative LMS-based DFE is then run on a subset of the inputs
 - Computes DFE for 100 samples
 - Time-step dependent!
 - Uses the CDR results above to get the digitized 1/0 waveform over which the LMS optimizes
- It applies the final DFE correction to the whole signal
- It runs again the filter (alpha, beta) and CR to get the final jitter

StatEye v5.2 Clock Recovery For Pattern Recovery (extractsignal.py)

- Conclusion:
 - StatEye v5.2 allows the user to specify input filtering
 - It does have an hard-coded jitter filtering
 - Default is a high-pass filter
 - Parameters should be scaled vs. time step
 - A 100-sample DFE is applied to the input for the CR
 - Input Filtering is applied also to the DFE-compensated result
 - A further hard-coded filter is used to convert the final jitter into a recovered « clock »