

Signal quality measurement method

Piers Dawe

NVIDIA

Supporters

- Karl Bois NVIDIA
- Tom Palkert
- Ali Ghiasi Ghiasi Quantum LLC
- Karl Muth Broadcom
- Alexander Rysin NVIDIA

Abstract

- At 200G/lane, C2M and CR receivers are similar
- Host loss means the traditional CR measurement method would be too far from its KR roots; it can be replaced by the better C2M method
- Also, the C2M signal quality measurement method can take advantage of learnings from TDECQ
- https://iee802.org/3/dj/public/24_05/calvin_3dj_01b_2405.pdf shows the practicality of the C2M eye method with the CTLE, FFE, 1-DFE reference receiver
- Observations on eye and jitter, SNDR, CR/KR voltage swing and observation bandwidth
- See last slide for comments that may be affected by the strategy

1. Method

Introduction

- Optical specs and C2M specs qualify **observable signals at defined interfaces with regard to their effect** on a reference receiver after the downstream part of the channel
 - Optical: TECQ and TDECQ
 - C2M: near end eye and far end eye
- For C2M, there is no separable or separately specified cable
 - One party is responsible for all the host; one party is responsible for all the module
- For optical, the defined optical channel is included or emulated
 - One party is responsible for all the transmitting host; one party is responsible for the optical channel; one party is responsible for all the receiving host

Specify what you care about

- Holistic quality metrics combine the effects of multiple causes of impairment without the expense and inefficiency of trying to diagnose each cause separately
 - TDP since 2002, then TECQ replaced a line-speed error-counting reference receiver with a more accurate scope-based receiver, then TDECQ added equalisation and PAM4
 - Eye mask and TDP proposal to replace jitter bathtub
 - https://iee802.org/3/ae/public/jan02/dawe_1_0102.pdf#page=17
 - Improved 100GBASE-SR4 transmitter testing
 - https://iee802.org/3/bm/public/mmfadhoc/meetings/may8_14/Improved100GBASE-SR4txTestingV1.pdf
 - Transmitter and SRS test source metrics
 - https://iee802.org/3/bs/public/15_09/king_3bs_01_0915.pdf
 - PAM4 TDEC straw-man conceptual proposal
 - https://iee802.org/3/bs/public/adhoc/smf/15_12_01/king_02a_1215_smf.pdf
 - Coherent (optical) specs are going in the same direction, looking **forward not backward**, towards specifying **results not causes**
 - https://iee802.org/3/dj/public/24_05/maniloff_3dj_02_2405.pdf

C2M

- Similarly, C2M assesses the signal at the compliance point with a reference receiver rather than deconstructing what might have caused it
- In 802.3ba (2010), C2M CAUI-4, the signal was assessed by eye width, eye height and vertical eye closure
- In Annex 120E (802.3bs, 2017), 50G/lane C2M, the signal was assessed by eye symmetry mask width, eye height and vertical eye closure
- In 802.3ck (2022), 100G/lane and P802.3dj D1.0, 200G/lane, the signal is assessed by eye height and vertical eye closure VEC
- A software channel is included in the scope as necessary (near end vs. far end, like TECQ vs. TDECQ for MMF)

CR history: increasing loss to the observation point

- In the KR method from 802.3ap (2007) the very low loss (back then) test boards allowed nearly direct measurements electrically near the package balls
- A backplane channel (linear with crosstalk) including its connectors was informative in 802.3ap (Annex 69A), assessed by
 - characteristic impedance, differential skew, IL(f), fitted IL(f), ILD(f), RL, ICR(f)
- Later backplane specs use COM as a holistic measure of channel suitability
 - COM assumes worst Tx and Rx ICs and packages
 - COM combines transmitter and receiver noise, crosstalk and jitter with channel properties
- A backplane spec divides the link into three sections of responsibility, but the boundaries are the package balls (now die bumps) not the connectors. The observable eye from the Tx to be qualified is not assessed directly for quality, but deconstructed for diagnosis. A stressful eye is used for receiver testing
- In 802.3ba (2010), the KR method was re-used for 40GBASE-CR4 and 100GBASE-CR10. Now the observation point was 6.5 dB from the package balls (85.8.3.4) and a de-embedding method was used
 - In CR, the IC in the host is required to be good enough for any recommended host channel (not just the one it's in) and the channel is recommended to be good enough for any KR-compliant IC (not just the one it is soldered to)
- In 802.3bj (2014), 100GBASE-CR4, the observation point was 9.85 dB away (92.8.3.6).
- In 802.3cd (2018), 50G/lane CR, the observation point was 10.07 dB away (136.9.3.2)
- In P802.3dj D1.0, CR, the observation point is expected to be up to $16.5+1.9+3.9 = 22.3$ dB from the die bumps (136.9.3.2). This is more than half way along the 40 dB channel
 - And very similar to the loss in 3ck C2M

What to do about it

- De-embedding over 20 dB, backwards towards the source, when we need to look forward to the point of use to find if the signal will be suitable for the reference receiver, less than 20 dB downstream, is misguided
- Diagnosing the de-embedded source and part-channel separately rather than assessing the observable signal at the compliance point, is misguided
 - R_{peak} and SNR_{ISI} for the channel, not well aligned to the reference receiver
- Salami-slicing the transmitter impairments into several different inflexible quotas is wasteful, and impractical for some impairments
 - Jitter, R_{LM} , SNDR. At least SNDR allows a trade-off between its two parts, and its measurement issue can be fixed
- It's time to modernise
- Measure the effect not the causes
- CR should use the C2M method
 - To align with C2M
 - Because the high loss between silicon and observation point makes the KR method impractical
 - Because salami-slicing the impairments is wasteful and "leaves margin on the table"
- KR and C2C should use the C2M method
 - To align with CR
 - Because salami-slicing the impairments is wasteful and "leaves margin on the table"

Consistent reference receiver

- The reference receiver in the scope and the COM reference receiver should be very closely aligned
- For CR,
 - High loss host is assessed with a software channel representing the remaining 17.7 dB loss allowed, and a reference receiver (same as the COM receiver) in a scope
 - Medium loss host is assessed with a 22.7 dB software channel and ...
 - Low loss host is assessed with a 27.7 dB software channel and ...
- The eye the scope calculates is very similar in all three cases
- For C2M,
 - For host output, the software channel is omitted
 - For module output, there is a software channel
- Reduced or omitted software channels are used ("near end eyes") as needed

Combine the quotas as FOM does

- In today's CR, a transmitter may trade off its voltage noise vs. its nonlinear distortion because they are both components of SNDR, but not its noise vs. jitter, v_f vs. R_{LM} , R_{peak} vs. SNDR... This is wasteful

Item	Combined in COM's FOM?	Combine in eye measurement?
<i>Steady-state voltage v_f (max)</i>	<i>Probably not needed if we have a V_{pkpk} spec</i>	
<i>Pulse peak ratio $R_{peak} = v_{peak}/v_f$</i>	<i>$\sim C_{eq} \sim EQ$ range. We don't yet know if we need fine TxFIR setting or not</i>	
Level separation mismatch ratio R_{LM}	No	Yes
SNDR part 1, noise	Yes	Yes
SNDR part2, distortion (but not R_{LM})	Yes	Yes
SNR_ISI	Yes	Yes
Jitter:		
J_{RMS}	Yes	Yes
J_{3u} , J_{3u_03} (J_{6u} has been proposed)	Yes	<i>Could be useful if it can be measured</i>
Even-odd jitter	No?	Yes?

Signal state of emphasis, inputs

- At present we believe that C2M will have more than two output modes. The reference receiver has evolved and the appropriateness of its reaction to state of emphasis is under discussion. We do not know if CR or C2M will need many finely spaced output modes as for the previous DFE-based receiver, or a handful. We do not know if "Tx FIR tap step size" is important or not
- Decisions on how to assess state of emphasis (rather than its quality and usability) are secondary and should follow the product signal specification strategy not usurp it
- C2M has stressed eye specs (e.g. 120G.3.3.5, 120G.3.4.3) with lossy electrical channel as appropriate
- CR has receiver interference tolerance and receiver jitter tolerance specs with lossy electrical channel as appropriate (low loss and high loss "tests")
- In both cases, calibration is quite involved and more precise than product signal specification. CR uses COM here as if it were a metric for a signal
- Decisions on whether noise loading or jitter loading or both, separately or combined, is the right choice, are independent of the product signal specification strategy

2. Improving C2M; jitter; SNDR

A good eye measurement recognises jitter

- TECQ and TDECQ are sensitive to a narrow eye caused by any combination of ISI, noise and jitter
 - They assess signals in relation to the target SER
- At 25G and 50G, C2M had an eye width spec
- https://iee802.org/3/ck/public/20_10/healey_3ck_01a_1020.pdf proposed two offsets (like TDECQ) for eye height and VEC
- After tens of dB of loss, marginal eyes end up looking much the same shape
- The difference between ESMW and two histograms is much like the difference between measuring width at a given height or height at a given width: both are set by the spec so it's about how we prefer to think about the measurement and talk about the margin
- Separate jitter measurements are remarkably troublesome. Partly because one is trying to measure one noise (the jitter) in the presence of other similar sized noises, from the signal under test and the receiver

SNDR dependence on observation loss is not as expected

- The way SNDR is defined and its relation to SNR_TX would need fixing
 - There is no need for separate SNDR and SNR_ISI measurements, at least in CR and C2M where we specify observable signals at compliance points
- If SNDR is kept as a separate item to be measured, it can be fixed
 - https://ieee802.org/3/dj/public/24_05/mellitz_3dj_02_2405.pdf
 - Comments 27 450 31 513 45 514 47
 - Comment 226 on SNR_ISI

3. Trend to lower voltages

- In 802.3ap (2007), max differential peak-to-peak voltage was 1600 mV for 1000BASE-KX, 1200 mV for 10GBASE-KX4 and 10GBASE-KR
- For C2M, the max is 900 mV except 3ck host output (750 mV)
- In the last 17 years, supply voltages and good-practice voltage swings have both come down
- Nearly every signal impairment is independent of signal swing
 - Except receiver nonlinearity
 - In COM, only η_0 implies what is enough signal, and COM's assumption that η_0 is independent of the signal is unrealistic
- For CR and KR, the max should be reduced from 1200 mV to 1000 mV
 - The range of v_f should be reduced similarly, from 400-600 mV to 350-500 mV

Comments 504 434 72 71 267 406 38 70 417

4. Observation bandwidth

- The observation bandwidth for measurements
 - Comments 422 425 141 432 131 133 137 439 230 60 399 245 32 309 410 124 225 412 217 415 53
 - Presumably it should be the same for C2M, CR, KR and C2C
 - Screening for ERL may need more bandwidth than measuring a signal, but it's done with a network analyser not a scope
- and the bandwidth in the reference receiver filter
 - Comments 257 404 36
 - Presumably this should be the same for C2M, CR, KR and C2C too, as the losses are more similar than in previous generations and the reference receiver architectures are expected to be similar
- and the filter type (Bessel-Thomson vs. Butterworth)
 - Comment 131
- are under discussion
 - We would like a high response in the pass band, low response from 67 GHz up, and benign phase response
- If we stay with a Bessel-Thomson filter, consider a 5th order filter which is the same at low frequencies, has better phase response, slightly better roll-off, and probably is more implementable, than 4th order
- At the same time, we should consider if the fixed third pole in the CTLE, which pre-dates the main receiver filter and should be seen as part of it, is justified

For reference: different metrics, electrical and optical equivalents

Electrical

Steady-state voltage v_f , VMA

Eye height

- Estimate of the net useful signal

VEC

- Imperfection of signal after reference receiver

(SNDR with SNR_ISI? and jitter)

- Imperfection (penalty) of signal before reference receiver

Optical

OMA

OMA-TECQ, OMA-TDECQ

TECQ- $10\log(C_{eq})$, TDECQ- $10\log(C_{eq})$

TECQ, TDECQ

May relate to these comments

- Keeping C2M method 132, 136, 139
- Proposes changing C2M to KR method 186, 187, 188, 189, 203
- Proposes improving the C2M method 523
- Proposes improving coherent method with a more complete signal quality metric 384, 579, 580
 - Assuming this will be decided on its own merits – mentioned here as a good example
- Proposes changing CR to C2M method 513, 514, 515
- Proposes improving the CR method 45, 47, 204 (236 is for KR)
- Jitter measurement and limits (236) 513 514 204 515
- SNDR 27 450 31 513 45 514 47
- SNR_ISI 226
- Voltage swing 504 434 72 71 267 406 38 70 417
- COM jitter parameters 504 143 72 71 271 272 70
- COM margin: C2M, CR RITT and JTol 249 400 401
- SNR_TX 31 504 143 72 139 71 270 41
- Observation BW 422 425 141 432 131 133 137 439 230 60 399 245 32 309 410
124 225 412 217 415 53
- COM f_r 257 404 36

Conclusion

- Build on the success of C2M
 - Learn from TDECQ
- High host loss in CR strains the KR method; we can take this opportunity to improve CR significantly
 - Simpler
 - Better use of available performance
- Reduce the signal swing for CR, KR and C2C in line with modern silicon, closer to C2M
- Consider observation filter types carefully