Suggested changes to COM

Adee Ran, Intel Corp.

Support and contributions

Supporters:

- Charles Moore, Avago
- Rick Rabinovich, Alcatel-Lucent
- Ed Sayre, NESA
- Liav Ben-Artsi, Marvell

Thanks to **Rich Mellitz**, **Arash Farhood** and **Adam Healey** for significant discussions and feedback that contributed to this presentation. They do not necessarily support my suggested remedies.

Annex 93a in a nutshell

- Create basic time-domain pulse response
  - From measurements, package model, RX filter
- Search for linear equalization setting (grid search over FFE and CTLE)
  - Select sampling phase
  - Apply "DFE window", with limited coefficient range
  - Calculate available signal A_s and PDF of residual-ISI cursors
  - Select setting that maximizes $FOM = \frac{A_s^2}{var(ISI)+\left(A_s\sigma_{RJ}\right)^2+\sigma_r^2}$
- Use selected LE settings to create noise PDFs
  - Apply CTLE and FFE to FEXT, create PDFs of "worst phase"
  - Apply only CTLE to FEXT, create PDFs of "worst phase"
  - Convert jitter PDF to noise PDF using available signal
  - Gaussian PDF with same σ_r used during search
- Convolve all noise PDFs, and calculate the BER quantile A_n of the joint PDF (where probability to exceed A_n is only BER)
- COM is the ratio A_s/A_n expressed in dB

Rationale
(should be consensus by now)

- Expected minimum receiver capability
  - Such receivers can be implemented using well-known methods
- Justifiable calculation of “noise” and “signal”
  - COM represents “closed eye ratio” of the reference receiver
  - Expected to correlate with real performance, and enable prediction
- Meaningful for both chip and board designers
  - Intermediate results can be used for “debugging”
- Open and efficient algorithm

Should we change anything?

- Let's look at each step… Is it broken?
  - **Generation of time-domain responses**
  - Tuning of equalization
  - Choice and calculation of noises
  - Meaning of the metric
- Also
  - Are we happy with the results?

Time domain response generation

- Calculated over many channels contributed during 802.3ap, 802.3ba, and 802.3bj
- With latest addition of package model – method creates meaningful results
  - One issue remain – real impedance at DC; Addressed by comment #101 (benartsi_3bj_01_0513)
- Overall seems satisfactory
  - No other issues identified in comments on D2.0

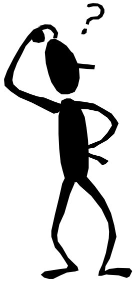

Should we change anything?

- Let's look at each step… Is it broken?
  - Generation of time-domain responses
  - **Tuning of equalization**
  - Choice and calculation of noises
  - Meaning of the metric
- Also
  - Are we happy with the results?

DFE limitation in KP4

- DFE limitation (for KP4) intended to prevent error propagation, but is applied to 1^{st} postcursor too
  - 1^{st} postcursor error propagation is mostly protected by the KP4 precoder; it is assumed that large value is allowed
  - This causes unjustified degradation for lossy channels in KP4
    - Example: Emerson longest: if change applied, 3.4 dB➔4.6 dB
- **Proposal:**

  Split the second case into two cases, n=1 and $2 \le n \le N_b$

  For n=1, use 1 instead of b_max

  For $2 \le n \le N_b$, use the existing equation

  - See backup for comparison

Comment #80; remedy also addresses #158 by Adam Healey

Noise in equalization tuning

- Noise taken into account in FOM (equation 93A-27) doesn't include crosstalk and DJ effect – unlike final COM calculation (equation 93A–38)
  - This was originally done to avoid repetitive calculation of crosstalk PDF for each equalization setting
  - Since total noise in FOM is smaller than in COM, the choice is always biased toward "stronger" LE; this generally results in degraded performance when crosstalk is added
  - Penalty is largest for "legacy" channels with significant crosstalk (may lead to unjustified failures)
- **Proposal: add crosstalk and DJ terms to FOM**
  - Avoid the unnecessary degradation.
  - Additional calculation burden isn't prohibitive.

Comment #74

Detailed suggested remedy to comment #74

- On each iteration of FOM calculation (93A.1.6)
  - In step a), add calculation of single bit responses of crosstalk sources, $h^{(k)}(t)$, using same procedure as thru
  - After step d), add a step that calculates the variance of each crosstalk source per-phase, as in equations 93A-36 and 37 (possibly moving them earlier in the text):

$$\sigma^{(k)} = \max_m \sigma_m^{(k)} \ , \ \sigma_{XT}^2 = \sum_k \sigma^{2^{(k)}} .$$

- In equation 93A-27, Change the denominator to

$$\sigma_{ISI}^2 + \sigma_{XT}^2 + \left(A_S\left(A_{DD} + \sigma_{RJ}\right)\right)^2 + \sigma_r^2 .$$

Should we change anything?

- Let's look at each step… Is it broken?
  - Generation of time-domain responses
  - Tuning of equalization
  - **Choice and calculation of noises**
  - Meaning of the metric
- Also
  - Are we happy with the results?

Channel noise sources

- COM includes a noise term σ_r that was intended to account for environmental noise (e.g. thermal, EMI) effective at the receiver's input bandwidth f_r.
  - See mellitz_01_0712 slide 7: "*Combine the RMS with a fixed white noise source before the CTLE*"
- But in the current procedure, σ_r is not affected by the selected CTF (which should typically attenuate noise). This effectively penalizes strong CTF boost in FOM, and causes sub-optimal results in COM.
  - Example: for the channel in mellitz_bj_01_0313 (35 dB channel for KP4), FOM selects 12 dB boost, which should reduce AWGN by ~3 dB
  - Since AWGN is the most significant noise component in this case, avoiding this reduction yields COM=2.1 dB (fail) instead of 4.1 dB (pass)!
- **Proposal: In equations 93A–27, 93A–32, and 93A–42, use a CTF-adjusted version of σ_r instead**:

Comment #73

$$\sigma_{r,eff}^2 = \frac{\sigma_r^2}{f_r} \int_0^{f_r} \left|H_{ctf}(f)\right|^2 df$$

Channel noise sources

- Additional issue pointed out by Adam Healey:
  - *"… the effective spectral density for your input-referred noise source is (sigma_r^2)/fr. Since fr is defined to be 0.75*fb, the spectral density for 100GBASE-KP4 is 25.78125/13.59375 ~1.9x that of 100GBASE-KR4. Is there any justification for the ambient noise being higher in a –KP4 environment than a –KR4 environment?"*
- Short answer: No.
- **Proposal: change σ_r in Table 94-19 from 1 mV to 0.5 mV**

No comment submitted…

Receiver noise sources

- "Implementation penalty" – Sensitivity/linearity/quantization/distortion etc.
- Clock recovery operates with relatively closed eye, effectively increasing jitter
- Mismatched terminations*
  - Which combination is worst? Depends on case
- COM currently allocates 3 dB to the receiver
  - ~30% vertical eye opening at the (relatively high) BER target
  - Seems too low to cover all of the above (see backup for experiment)
- Other changes proposed here increase COM values of "limit" channels
  - Do we want tougher channels?

* Not receiver noise per se, but has an effect

Receiver noise sources

- **Proposal:**
  - **Change minimum from 3 dB to 4 dB for KR4, and to 5 dB for KP4**

Comments #103, #104

- **Alternative proposal made in ad hoc meeting:**
  - **Increase the values of σ_r instead, so that margin is reserved for long channels, and shorter channels are not penalized**
  - **Change from 1 mV to 1.6 mV for KR4**
  - **Change from 1 mV to 1.2 mV for KP4 (instead of making it 50% of KR4; this adds margin for KP4 receiver penalty)**

Not a part of submitted comment remedy…

Transmitter noise sources

- TX Jitter is limited and accounted for in COM – A_{DD} and σ_{RJ}
  - Current A_{DD} value used in clause 94 (0.05 peak) does not match TX CDJ spec (0.05 UI PTP)
  - **Proposal: change A_{DD} in Table 94-19 to 0.025 UI**

Comment #82

- Additional noises include package crosstalk, power supply noise, etc.
  - Limited by specifications, but differently for KR4/CR4 (as normalized linear fit error – RMS across phases) and KP4 (as SNDR – RMS per phase)
  - It is assumed that the specified TX noise limits are low enough to be tolerable by receivers.
    - Not simple to analyze in general
    - Validation depends on receiver implementation – and is beyond the scope of the standard.
  - **No change proposed at this point**

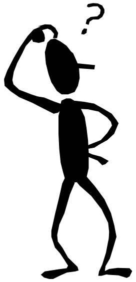

Calculation of noise PDFs

- Current recommendation is to use voltage bin size of 0.1 mV.
- For long channels, available signal (after equalization) is only a few mV, and crosstalk responses might be isolated spikes of less than 0.1 mV amplitude
  - Recommended bin size causes ignoring most of the crosstalk effect
- Reducing bin size to 10 µV accounts for most crosstalk
  - In one case analyzed, impact is ~1 dB (see backup)
  - Below 10 µV: diminishing returns
- **Proposal: recommend max bin size of 10 µV**

No comment submitted…

Should we change anything?

- Let's look at each step… Is it broken?
  - Generation of time-domain responses
  - Tuning of equalization
  - Choice and calculation of noises
  - **Meaning of the metric**
- Also
  - Are we happy with the results?

What does COM mean?

- "COM is practically the relative vertical eye closing at the BER probability, expressed in dB"
  - COM defines a reference receiver; a given receiver may perform better – but COM was introduced for channel qualification, not for comparing receivers
  - Real receivers performance does depend on eye opening/closing
- COM is not the classic "receiver margin"
  - Receiver margin means "how much can we increase X and still work" where X can be: additive noise, transmitter jitter, crosstalk…
  - **These are receiver metrics – not channel metrics**
  - If we are worried about unknown stresses – let's make them known!
- **No change proposed**

Should we change anything?

- Let's look at each step… Is it broken?
  - Generation of time-domain responses
  - Tuning of equalization
  - Choice and calculation of noises
  - Meaning of the metric
- Also
  - **Are we happy with the results?**

Are we happy with the results?

- Current program shows failures for several channels, mainly at 100GBASE-KP4 signaling
- With the proposed changes, COM values are generally increased; “maximum loss” channel for KR4 yields COM ≈ 4 dB, and for KP4 yields COM ≈ 5 dB

Sample result comparison

Channel	PMD type	COM (version post-D1.3a)	COM (proposal)
35db Loss (IBM)	100GBASE-KR4	4.4 dB	4.9 dB
mellitz_bj_01_0313	100GBASE-KP4	-0.5 dB	4.34 dB
Emerson longest	100GBASE-KP4	3.4 dB	4.8 dB
TE Whisper 42.8” Meg6	100GBASE-KR4	8.1 dB	8.1 dB

Note: “proposal” results are based on remedies of my original comments. Other comments and proposals that may affect COM have been submitted; their effect is not included.

Backup

DFE limitation in KP4 – illustrated

Current procedure

- Resulting available signal: 14 mV (44/3)
- 1st postcursor: 0.19*cursor
- COM=3.45 dB

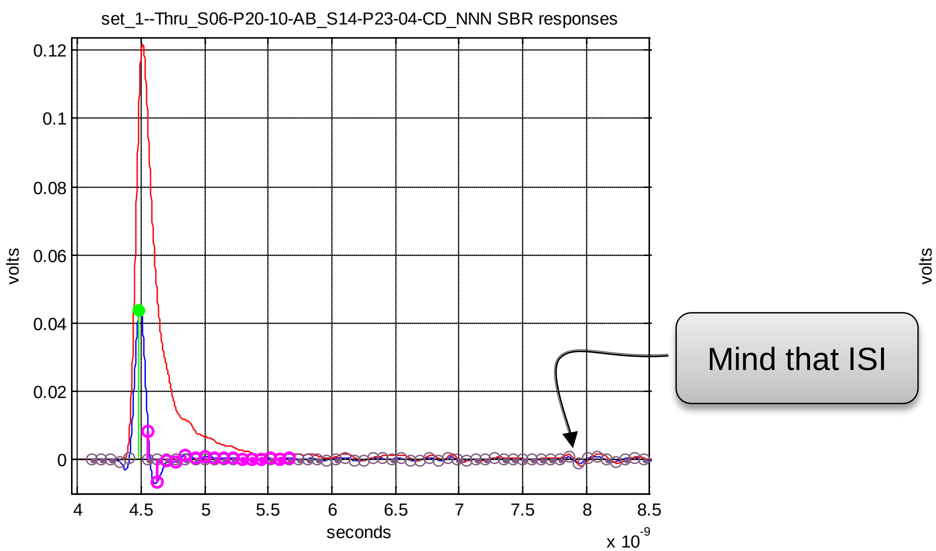

Proposal

- Resulting available signal: 19 mV (57/3)
- 1st postcursor: 0.55*cursor
- COM=4.49 dB

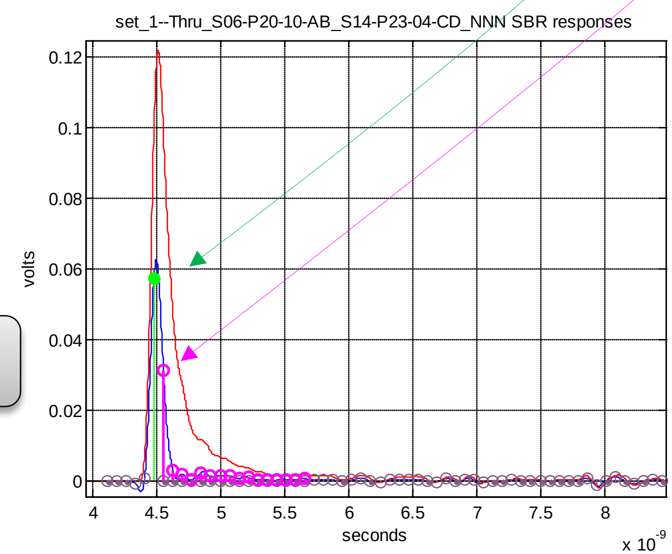

Channel used: Emerson longest, Thru_S06-P20-10-AB_S14-P23-04-CD_NNN, 1 FEXT and 4 NEXT

COM effect on receiver jitter

- 3 dB means that at the target BER, the vertical eye is 10^(-3/20)=70% closed
- Practical receiver implementation can only make it worse (quantization etc.)
  - But say the actual eye is still not completely closed – is it OK?
- The receiver is required to track low frequency jitter; required tracking capability ("bandwidth") is implied from TX jitter specs, and verified in RX jitter tolerance tests
  - Closing the eye reduces the effectiveness of phase detectors
  - For a given "bandwidth", the more closed the eye, the more receiver jitter is created.
- Receiver jitter is an "implementation penalty" that depends on COM!

Budget consumption experiment

- 35db_Loss_Channel (IBM), NRZ
  - Baseline after proposed changes: 4.91 dB
  - Refine PDF quantization (bin size) from 1e-4 to 1e-5
    - Beyond that – diminishing returns
  - Add DFE quantization
    - Step size: 5% of cursor
    - Can probably be better
  - Add RX jitter
    - E.g. due to noisy CDR
    - Assume mostly DJ; take 2*A_{DD} instead of A_{DD}
  - Add worst case impedance mismatch
    - In this case TX: 45 Ω, RX:55 Ω

Results…

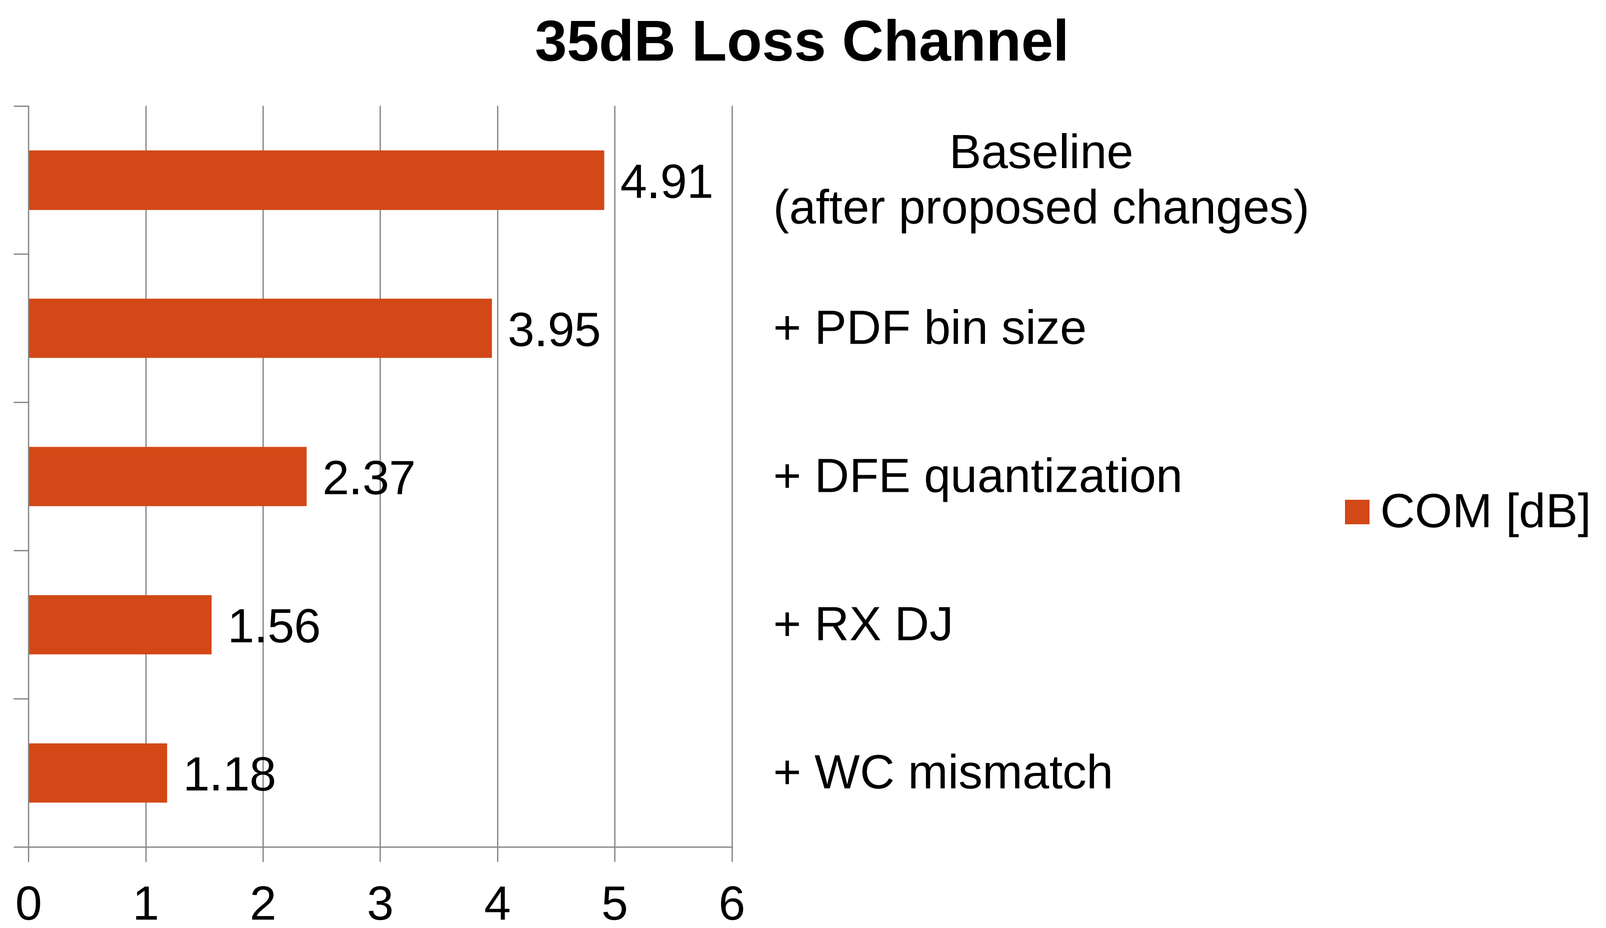
35dB Loss Channel
Baseline
(after proposed changes)
+ PDF bin size
+ DFE quantization
+ RX DJ
+ WC mismatch
4.91
3.95
2.37
1.56
1.18
0
1
2
3
4
5
6
COM [dB]

Conversion of jitter to noise

- Simplifying assumption: jitter moves the sampling phase along the slope of the rising edge.
  - The pulse starts at a zero crossing at the precursor and reaches the signal level A_S at the cursor
  - A linear approximation is that the slope is A_S [volts/UI]
- Thus, jitter is converted to voltage noise by linear scaling of the jitter PDF to noise PDF, and the conversion factor is A_S – easy to specify and compute.
- Precursors and postcursors typically have smaller slopes so their “contributions” are lower. Accounting for them makes the analysis much more complex, without significant change of the result.
  - Requires calculation of conditional probabilities of cancelling/adding-up combinations of slopes.
  - Not as easy as convolutions…

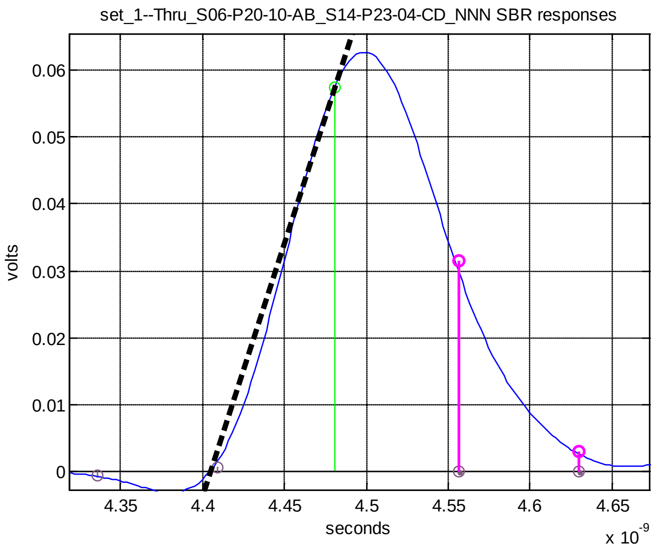

PAM-4 jitter effect

- For PAM4, the transitions are not only between two levels; but we analyze the transition between two adjacent levels as the worst case (though the ISI and crosstalk statistics include all possible levels and transitions).
  - In this case, the effective signal level is 1/3 of the amplitude used in NRZ. Other transitions indeed suffer from larger jitter effect, but the rest of the noise sources do not scale, so the "worst case" is still valid. The fact that not all transitions are worst case does leave a small margin in the COM result for PAM4.
  - A more correct calculation would consider all possible transitions:

$$p(error) = \sum_{\Delta \in transitions} p(error|\Delta)p(\Delta)$$

  - 2/3 transitions are the most frequent type (6 out of 16), so they have the maximum effect on the overall DER (detector error ratio).