Initial ILT setting and maximum Tx swing

Comments #218, #263

Adee Ran, Cisco

The comments

Cl **179** *SC* **179.9.4.1.3** *P* **383** *L* **31** # 263

Ran, Adee Cisco

Comment Type **TR** *Comment Status* **D** *Max swing & initial ILT setting*

The "initialize" values adopted in D1.4 are different for CR and for C2M.

This requires different initialization in the transmitter and, very likely, a different algorithm in the receiver, depending on the mode chosen for the port (whether a module or a copper cable is plugged). These create an unnecessary burden for firmware developers, possibly increasing the code size and development/debugging time.

The motivation for choosing preset 6 for the initial setting was to limit the initial swing reaching the receiver input. The maximum transmitter swing with preset 6 is 0.75 V. In comparison, CR initial setting is preset 1, which has a maximum transmitter swing of 1 V.

It is reasonable to assume that CR receivers can handle 1 V output swing of the transmitter (which will be attenuated by the channel, assumed to have considerable loss at frequencies present in the ILT signal).

If preset 6 is used as the initial value for CR too, the transmitter's v_f (measured near the transmitter with preset 1) for these PMDs can be allowed to be as high as 0.6 V; If a device has v_f at this maximum value, then with preset 6, the transmitter swing will be 0.9 V, lower than the 1 V currently allowed. If a device has v_f of 0.5 (the maximum in D1.4) its maximum will be 0.8 V. Either way, the receiver will see an even lower swing.

This will enable using a higher output swing for CR, potentially increasing their reach (if the transmitter is capable), and using the same adaptation algorithms in the receiver.

This change does not require increasing A_ne in COM; having transmitter swing at the maximum on one end of the cable and at the minimum on the other is not a likely situation and can be excluded from cable compliance assumptions. Devices should work with cables that meet the existing specifications.

A similar argument can be made for KR vs. C2C.

SuggestedRemedy

In Table 179-7, change the Transmitter steady-state voltage v_f range from "0.4 to 0.5" to "0.4 to 0.6", and change "differential peak-to-peak voltage (max) , transmitter enabled" from "1" to "1.2".

In Table 179–8, change the "initialize" setting to match preset 6, and delete "and initialize" in the footnote.

In Table 179–10, change the "Amplitude tolerance" value from "0.5" to "0.6".

in 179.9.5.2, add an informative note as follows:
"NOTE--The steady-state voltage in Table 179-10 corresponds to preset 1. It is not initially generated by a transmitter, due to the initialize setting in Table 179–8. The receiver is not required to tolerate preset 1 unless it specifically requests for it."

Optionally, apply the corresponding changes in clause 178.

Cl **179** *SC* **179.9.4.1.3** *P* **383** *L* **31** # 218

Dawe, Piers Nvidia

Comment Type **TR** *Comment Status* **D** *Max swing & initial ILT setting*

Transmitters are supposed to start Training at medium amplitude (preset 6) now, not the loudest, to avoid possible crosstalk and linearity issues. A receiver that prefers a louder signal on a particular channel can ask for it.

SuggestedRemedy

In Table 179-8, for "initialize", change 1 to 0.75, add the tolerances, and delete "and initialize" in the table footnotes. As in Table 176D-9 (which applies to 176C).

Proposed Response *Response Status* **W**

PROPOSED ACCEPT IN PRINCIPLE.
For CRG discussion, along with comment #263.

The suggested changes

Both comments suggest

A. Changing the "initialize" setting for CR PHYs to have c(0) of 0.75 (as in preset 6) instead of 1 (as in preset 1), identical to AUI-C2M.

(for different reasons: #263 is about transmitter initialization logic and receiver adaptation algorithms, #218 is about crosstalk and linearity).

#263 additionally suggests

B. Adding an informative note explaining amplitude tolerance (179.9.5.2), as in the proposed response to comment #126 (which addresses amplitude tolerance in C2M).
C. Changing the maximum steady-state voltage for CR PHYs from 0.5 V to 0.6 V and make corresponding changes in the peak-to-peak (1.2 V) and amplitude tolerance specifications.
D. Considering change C for clause 178 too.

This presentation assumes that item A (common to both comments) and item B (non-controversial) will be implemented; it focuses item C.

Why is max steady-state voltage 0.5 V?

- The maximum v_f was 0.6 V (1.2 Vpp) in D1.2, and in many previous electrical PMD clauses.
- The change to 0.5 V was proposed in simms_3dj_01_2411 with reasoning shown in the next slide.
  - Straw poll #E-1 indicated support: Y: 22 N: 10 NMI: 5 A: 9 (see response to comment #345 against D1.2)
  - The proposal was adopted and implemented D1.3.
- In D1.4 there was an additional change of the initial equalizer value in the AUI C2M/C2C transmitter…
  - Further reducing the signal at the receiver input
  - This change was not applied for CR/KR PMD transmitters

The reasons for the reduction

Summary

Reducing Vf max from 0.6V to 0.5V

- Reduces complexity of receiver design
- Enables lower power supply SERDES
- Enables current and future silicon process nodes
- Increases compatibility with other SERDES Specifications
- Reduces system power and noise

Source: simms_3dj_01_2411 slide 9

- The main technical argument was the complexity of receiver design to handle a large input swing, especially with current and future silicon process nodes.
  - Enabling lower power supply is likely related to that.
- Other benefits mentioned seem less relevant – since they could be achieved by design choice even with the previous limits.
- It was also mentioned that COM improves by this reduction.

Addressing receiver input swing

- The initial signal sent over the cable is training frames at the “initialize” setting, synchronous PRBS13 and PAM2 modulation.
- Even if a transmitter has the maximum allowed swing, the signal seen by the receiver will be attenuated by the channel
  - And by the “initialize” setting, assuming we adopt the proposed change (A).
- For CR, there is a minimum cable assembly loss, specified as 16 dB (Table 179–13)
  - This is between TP2-TP4
  - TP0d to TP5d is likely significantly higher!
- A receiver can (and likely will) change the Tx equalization setting through ILT; for a high loss channel, it is likely that the full swing (maybe with some Tx equalization) will be beneficial
  - What setting is used depends on the receiver’s ILT algorithms
- Simulation results in the subsequent slides show the effect of a 16 dB channel and Tx equalization on the Rx input swing.

Simulation setup

- Pattern: two training frames with PRBS13 PAM2,
- Tx output swing: 1 Vpp
- Test cases:
    1. 1 dB channel (very close to Tx output) with preset 1 (initial training pattern in CR)
    2. 16 dB channel (Rx input, minimum IL) with preset 1
    3. 16 dB (Rx input, minimum IL) + preset 4

Note: preset 4 (nominal coefficients [0.05, -0.2, 0.75]) creates a DC voltage attenuation of 0.6 (-4.4 dB).

16 dB Channel used in the simulation
(smooth channel created using the PCB mathematical model)

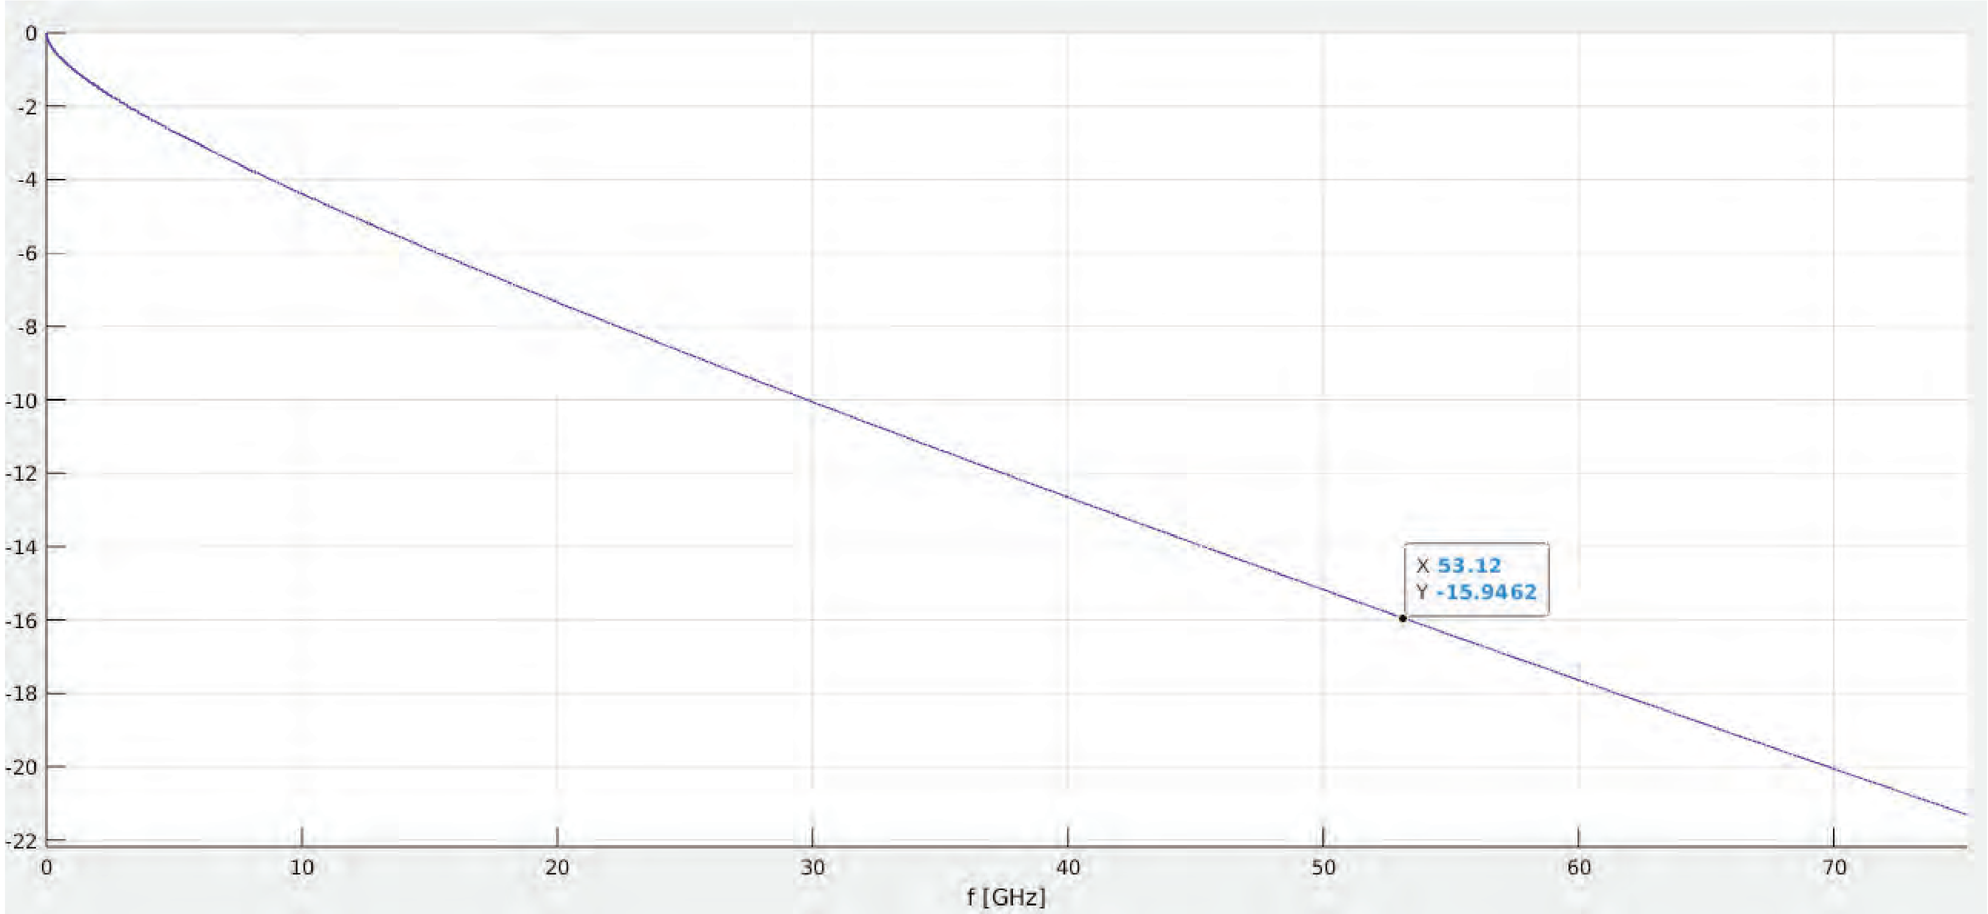

Waveforms: two full training frames

Marker and DME region

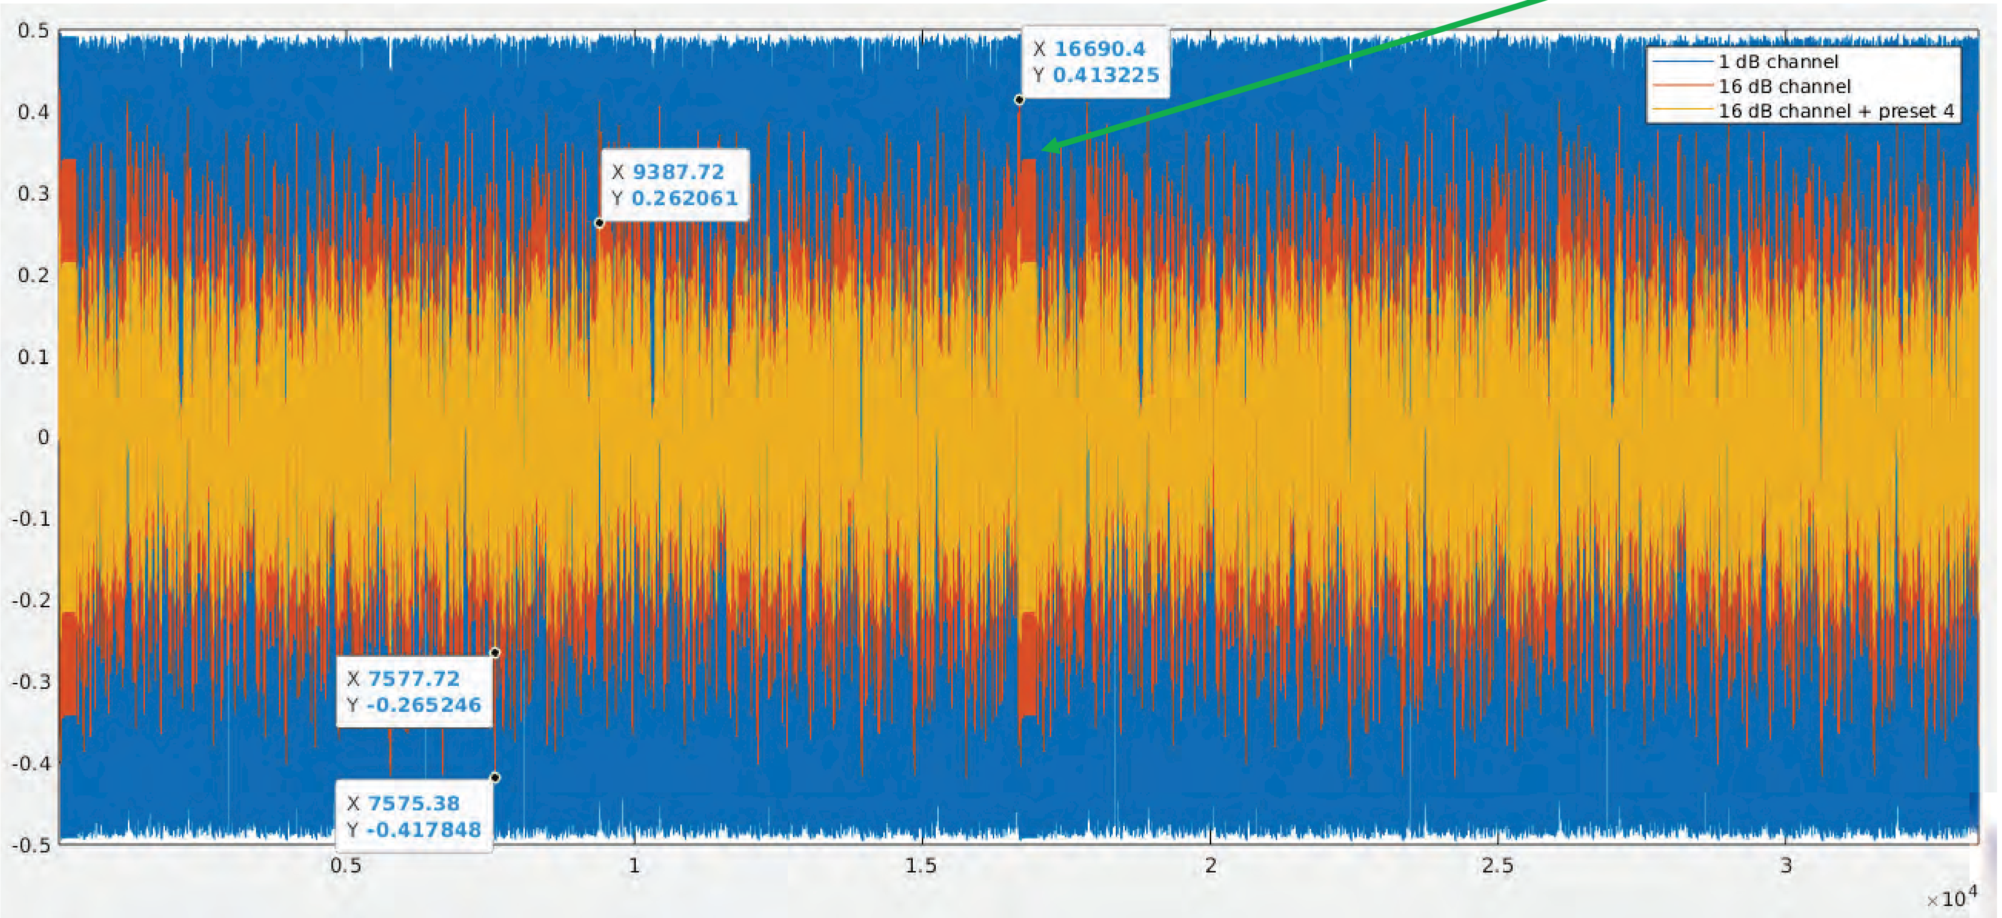

The signal is attenuated by the channel and by the Tx equalizer – the Tx output swing is never seen by the receiver

Zooming on the training frame marker + DME

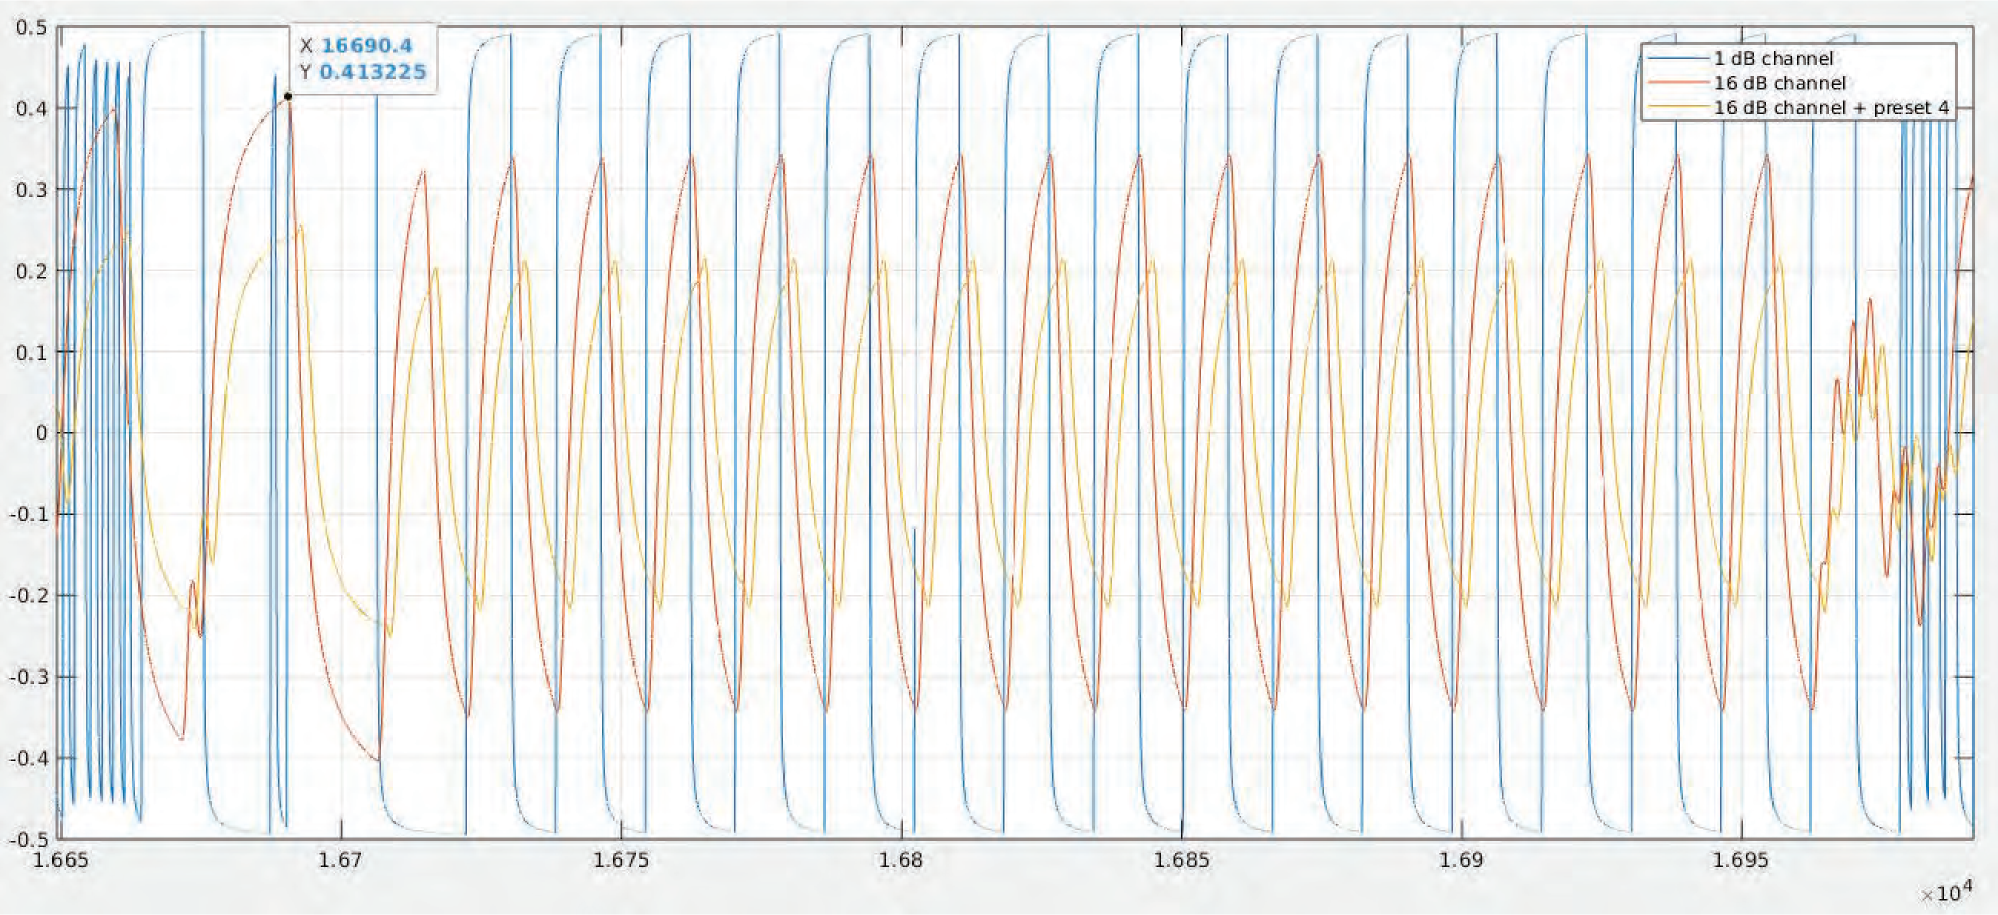

Observations

- The 16 dB channel reduces the swing of the training signal from 1 Vpp (Tx) to **0.83 Vpp** (Rx)
  - If preset 6 is used for initialize (change A), it would become **0.62 Vpp** at the Rx input
  - If preset 6 is used for initialize *and* maximum Tx swing is increased to 1.2 Vpp (change C), it would become **0.747 Vpp** at the Rx input – still lower than the current 0.83 V
- Applying preset 4 in the Tx reduces the Rx input swing to **0.53 Vpp**
  - If maximum Tx swing is increased to 1.2 Vpp (change C), it would become **0.636 Vpp** at the Rx input
  - With preset 4, the maximum possible Rx input swing in data mode is **0.72 V** (extremely rare long runs)
- **The input swings listed above should not be a problem for current Rx designs**
  - Experience from 100G SerDes confirms that.
- Practical CR channels have loss much higher than 16 dB – so the values above are conservative.

Summary of results

Specification	Initial Tx output swing [V]		Initial Rx input swing, 16 dB channel, PRBS13 [V]		Max Rx input swing, 16 dB channel, data mode with preset 4 [V]	
	Min	Max	Min	Max	Min	Max
802.3ck PMDs (e.g. 200GBASE-CR2)	0.8	1.2	0.73	1.09	0.8*0.6=0.48	1.2*0.6=0.72
200GBASE-CR1 (D1.3 and current) ***200GAUI-1 in D1.3***	0.8	1	0.7	0.87	0.8*0.6=0.48	0.6
200GAUI-1 (current)	0.8*0.75=0.6	0.75			0.8*0.6=0.48	0.6
200GBASE-CR1, per comment #263 (A)	0.8*0.75=0.6	0.75	0.52	0.65	0.8*0.6=0.48	0.6
200GBASE-CR1, per comment #218 (A&C)	0.8*0.75=0.6	1.2*0.75=0.9	0.52	0.78	0.8*0.6=0.48	1.2*0.6=0.72

What about COM?

- COM is a channel metric, not a predictor of real link performance
  - The combination of A_v, A_fe, A_ne (victim, FEXT, and NEXT amplitudes) has always been "the worst combination"
  - Thus, the values we have today are based on min v_f (0.4 V) and max v_f (0.5 V)
- What if we increase max v_f to 0.6 V?
- Assume some device uses that allowance and has v_f above 0.5 V
  - Its partner has a larger signal (and FEXT) – it is generally beneficial for Rx performance (if not, it can attenuate it)
  - The device has to endure the higher NEXT it creates on its own Rx
    - This may be within the implementation budget (3 dB COM)
    - It should be an implementation choice
- If both sides of a link have the same v_f, increasing it can provide benefits:
  - Longer reach for the same Rx performance
  - Better Rx performance (lower BER, FLR) for a given channel
- The PMD specification should allow that
- We can keep the COM amplitude parameters as they are in D1.4, or change them to match v_f
  - Either way, this can be done after D2.0
  - Note that cable assembly channels (TP2 to TP4) have not been contributed so far…

What about transmitters?

- The v_f range min/max should cover both design freedom and variations between devices.
- The reduced range limits design freedom.
- Given receivers can handle it, whether Tx implementations will reach the maximum swing should be an implementation choice
  - Perhaps no implementation will reach v_f of 0.6 for practical reasons…
  - But the current maximum will unnecessarily limit implementations.
- Enabling higher swing can be used to increase cable reach, which can reduce overall power by removing the need for retimers or reducing the number of hops in a network.
  - An unnecessary limiting spec might be ignored by the market!

What about clause 178 (KR)?

- The same arguments made for CR also apply to KR
  - Except that there is no defined minimum channel IL
- How likely is it that a backplane system has ILdd (TP0d to TP5d) lower than 16 dB?
  - This possibility does not seem to justify a different specification.
  - Tx output can be reduced by device configuration.

Summary

- Referring to the changes listed on slide 3:
  - **Change A** is suggested by both comments and will hopefully be in consensus.
  - **Change B** provides clarification of amplitude tolerance, and matches a proposed response to #126.
  - **Change C** would enable broader application space for CR PHYs, while not harming receivers, as shown.
  - **Change D** would align the KR and CR specifications.
- Recommendation: implement all four changes.

That’s all

Questions?