Error budgets for AUIs within a PHY

Adee Ran, Cisco

Background

- In ran_3dj_elec_02_230622 it was noted that:
  - A DER_0 value of 2.67e-5 was adopted for the higher-loss AUIs within a PHY (see motion #8 and ran_3dj_02_2305)
  - This corresponds to BER of 2e-5 with uncorrelated errors, or measured BER of 4e-5 with precoding ON
  - As part of that decision, **BER division between C2C and C2M** and measurement method were left to be determined
- This presentation addresses the division between C2C and C2M.
  - Any reference to “BER” or measurement method for future interfaces is only for illustration purposes.

Possible PHY structures (from 802.3df D2.1)

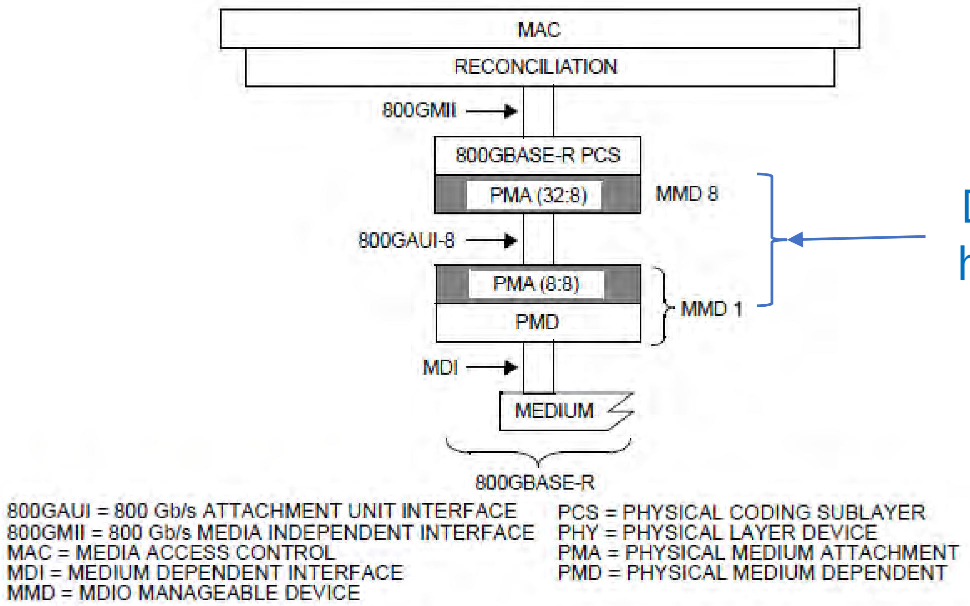

Figure 173A–1—Example PMA layering with eight-lane PMD and single 800GAUI-8

The adopted DER_0=2.67e-5 holds for both cases

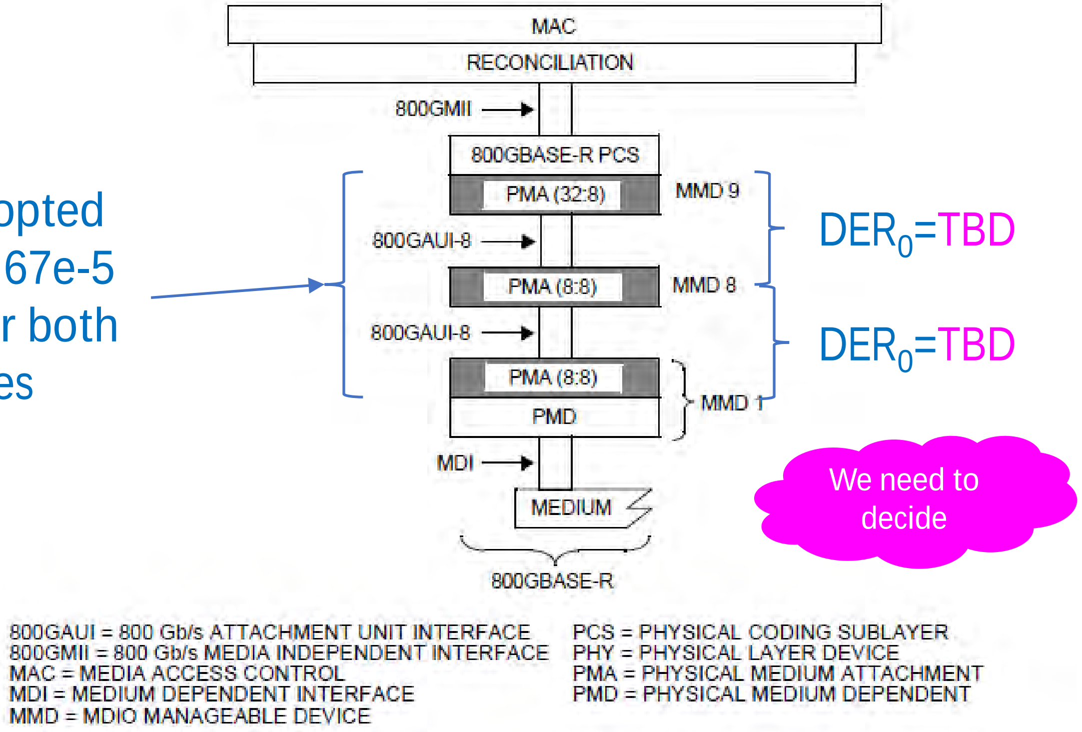

Figure 173A–2—Example PMA layering with eight-lane PMD and two 800GAUI-8

AUIs within an xGMII Extender are not within the PHY – not addressed in this presentation

BER allocations in existing AUI specifications

120E.4.2 Eye width and eye height measurement method

Eye diagrams in 200GAUI-4 and 400GAUI-8 chip-to-module are measured using a reference receiver. The reference receiver includes a fourth-order Bessel-Thomson low-pass filter response with 33 GHz 3 dB bandwidth, and a selectable continuous time linear equalizer (CTLE) to measure eye height and width. The pattern used for output eye diagram measurements is PRBS13Q. Unless specified otherwise the probabilities are relative to the number of PAM4 symbols measured. The following procedure should be used to obtain the eye height and eye width parameters, as illustrated by Figure 120E–13:

1) Capture the PRBS13Q using a clock recovery unit with a corner frequency of 4 MHz and slope of 20 dB/decade. The capture includes a minimum of 3 samples per symbol, or equivalent. Collect sufficient samples equivalent to at least 1.2 million PAM4 symbols to allow for construction of a normalized cumulative distribution function (CDF) to a probability of 10^{-5} without extrapolation.

120F.3.1.1 Peak-to-peak AC common-mode voltage

The low-frequency and full-band peak-to-peak AC common-mode voltage, VCM_{LF} and VCM_{FB}, respectively, are defined by the method specified in 162.9.4.4 with the exception that the peak-to-peak AC common-mode voltage is defined as the AC common-mode voltage range measured at TP0v that includes all but 10^{-5} of the measured distribution, from 0.000005 to 0.999995 of the cumulative distribution.

The low-frequency peak-to-peak AC common-mode voltages shall meet the specification for VCM_{LF} (max) in Table 120F–1.

Table 120F–5—Receiver interference tolerance parameters

Parameter	Test 1 (low loss)			Test 2 (high loss)			Units
	Min	Max	Target	Min	Max	Target	
FEC symbol error ratio[a]	—	10^{-4}	—	—	10^{-4}	—	—
Insertion loss at 26.5625 GHz[b]	9.5					—	dB
RSS_DFE4[c]	0.05	—	—	0.05	—	—	—
COM including effects of broadband noise[d]	—	—	3	—	—	3	dB

(Equivalent to BER $=10^{-5}$)

120G.3.3.5 Host stressed input tolerance

Host stressed input tolerance is defined by the procedure described in 120G.3.3.5.1 through 120G.3.3.5.3.

The host under test shall meet the BER requirement in 120G.1.1: (Less than 10^{-5})

- For either the short or long mode.
- With all sinusoidal jitter cases in Table 162–17.
- For any signaling rate in the range given in Table 120G–7.

120G.3.4.3 Module stressed input tolerance

Module stressed input tolerance is defined by the procedure described in 120G.3.4.3.1 through 120G.3.4.3.3.

The module under test shall meet the BER requirement in 120G.1.1: (Less than 10^{-5})

- For both the high-loss and low-loss cases (see 120G.3.4.3.2), calibrated separately.
- With all sinusoidal jitter cases in Table 162–17.
- For any signaling rate in the range given in Table 120G–9.

120G.5.1 Signal levels

The signal levels are as defined in 120E.3.1.2.

Low-frequency and full-band peak-to-peak AC common-mode voltage, VCM_{LF} and VCM_{FB}, respectively, are defined by the method specified in 162.9.4.4 with the following exceptions:

a) The peak-to-peak AC common-mode voltage is defined as the AC common-mode voltage range measured at TP1a or TP4 that includes all but 10^{-5} of the measured distribution, from 0.000005 to 0.999995 of the cumulative distribution.

b) The condition for transmitter equalization to be turned off does not apply.

BER budget division in the existing standard

- As shown, existing C2C and C2M are specified with a maximum BER of 1e-5 each
  - Modules have BER<1e-5 specified for stressed input test with specific test channels
  - Host output specified with EH and VEC with a probability of 1e-5
  - Similarly for module output and host input
  - Similarly for C2C
- The total PHY allocation within the RS-FEC error budget is 2e-5 per PHY
- What if a host has only C2M and no C2C?
  - Host output still needs to meet EH and VEC with a probability of 1e-5 – it does not get to use the total 2e-5 (so, for example, reach can't be extended)
  - Host stressed input has the same signal, and the requirement is still BER<1e-5 – it does not get a benefit
  - Module specifications do not change
  - The BER provided to the RS-FEC is lower than the maximum – leaving margin
- The case of a host that has only C2C and no C2M is similar
  - Possible with passive copper cable link

Possible host design with 200G/lane AUIs

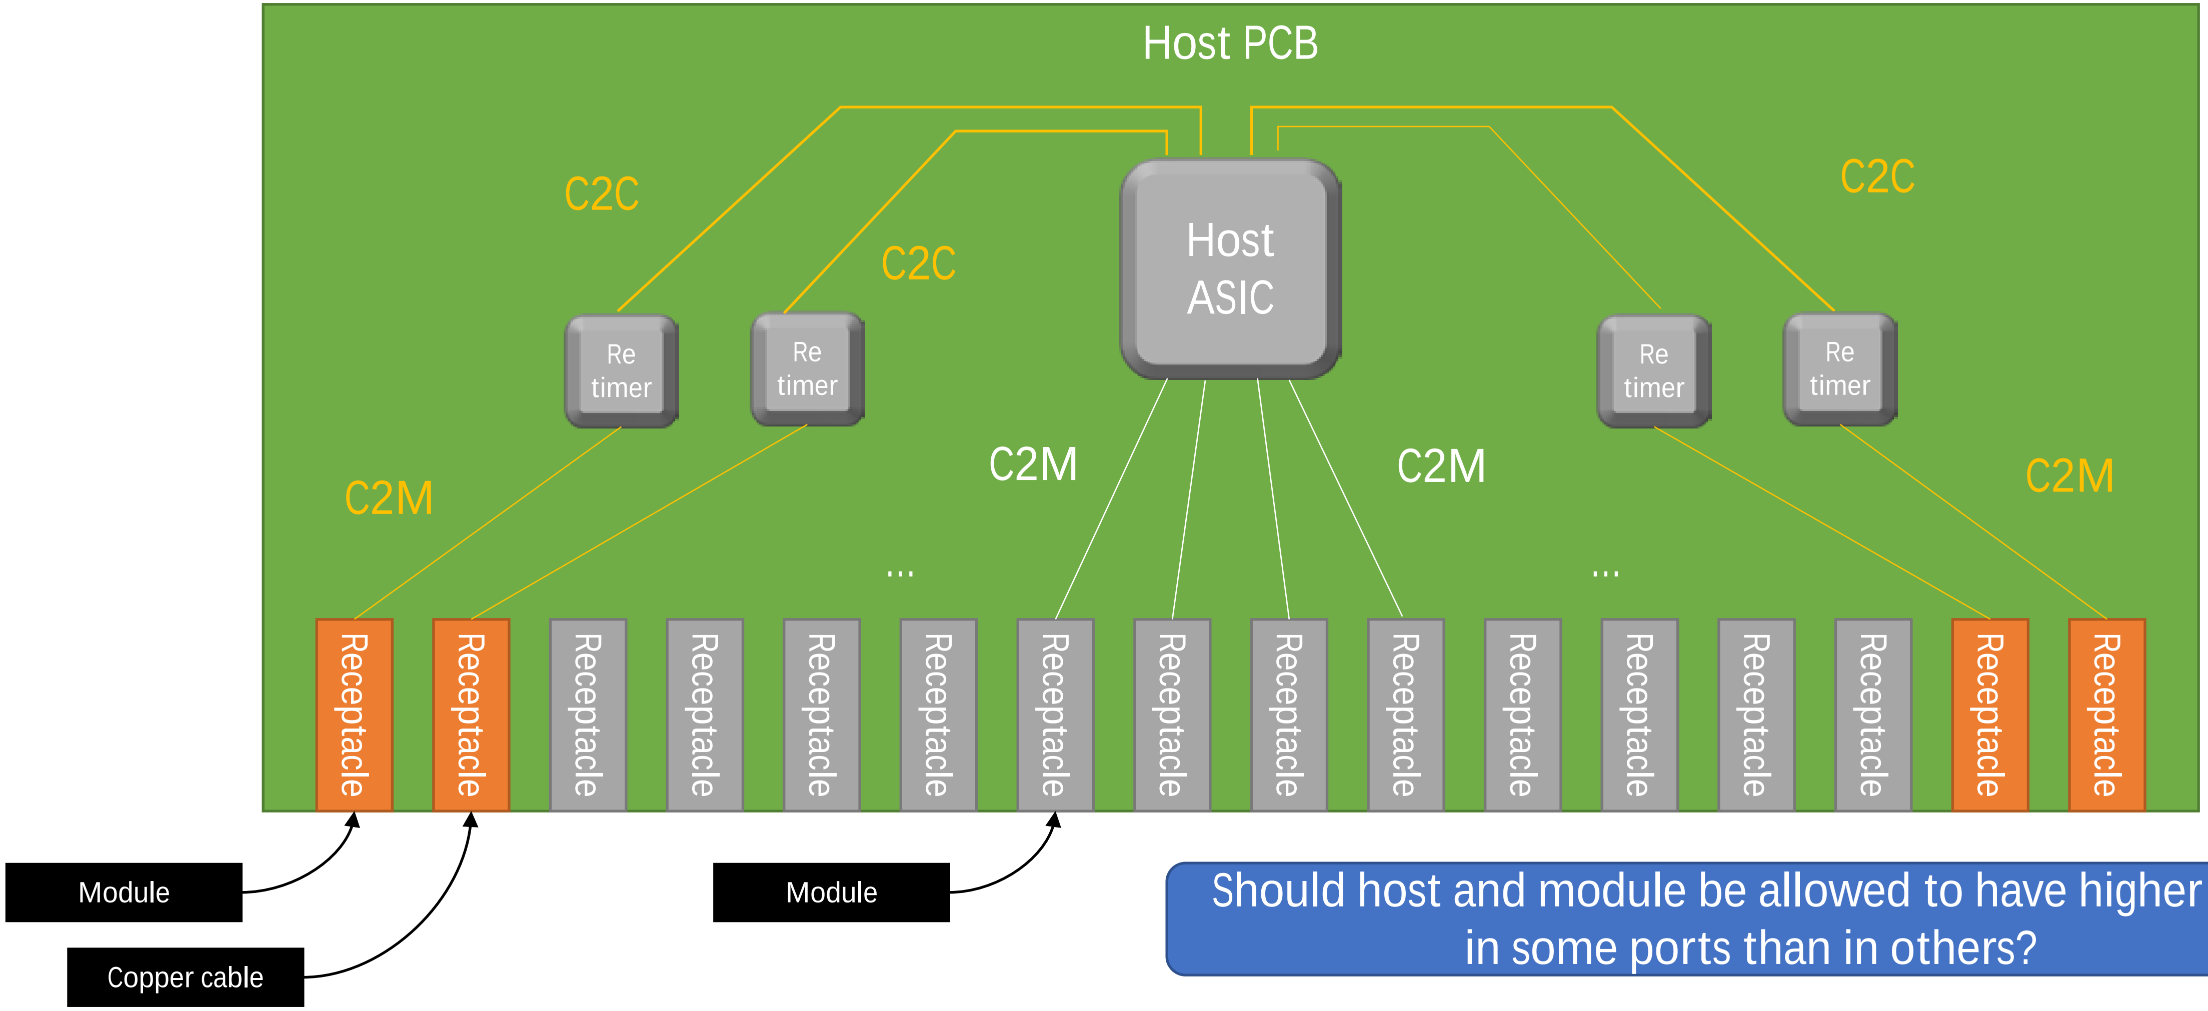

Should host and module be allowed to have higher BER in some ports than in others?

Flexible or fixed allocation?

- Approved motion text: “Adopt a DER0 value of 2.67e-5 (…) as the total allocation for higher-loss AUIs within a PHY”
  - This is consistent with “BER_{Module}<2.4e-4” and enables progress in optical specs!
- It we interpret the decision as “a PHY can always use the total allocation”, then:
  - A port with C2M and no C2C (optical module, no retimer) should have higher BER allocation for the C2M than a port with C2C+C2M (optical module with retimer)
    - Since BER is part of the specifications of both host and module, this mean two host specifications, e.g., stressed input calibration and required BER
    - And similarly, it will lead to two module specifications
  - Likewise, a port with C2C and no C2M (passive copper with retimer) should have higher BER allocation for the C2C
    - But if an optical module can be plugged into the same port – the budget should be split
    - Beneficial only for hosts that will only work with copper
- Alternatively, we can have fixed allocations for C2C and C2M, as in previous projects.
  - As mentioned in lit_3dj_01_2307, the impact is ~0.34 dB reduction in COM when only one is used.

Considerations for P802.3dj

- Many hosts do not use retimers
  - Effectively the C2C budget is not used; with fixed allocation, margin is left on the table
  - For such hosts, having 100% of the budget for the C2M would allow more flexibility in channel design
- Some hosts may be designed only for copper cable, and have retimers
  - For such hosts, having 100% of the budget for the C2C would allow more flexibility in channel design
- Can we have different C2M error allocation based on whether C2C is used or not (and vice versa)?
  - Hosts can choose one specification based on their structure
  - For modules to work with any kind of host – it would mean two sets of compliance specifications (could be identical except for the BER and its effect on measurements)
  - A host that never has C2M (copper only) can have higher BER for the C2C
- We could also give C2M 100% of the budget by removing the option to have C2C within optical PHYs
  - A C2C interface can still be used outside of the PHY if the external chip has a PHY XS+PCS. In that case it can have a much higher BER.
  - Modules will have one specification.
- Backward compatibility
  - In a system that has both C2C and C2M, and one of them is at 100G/lane, the budget for the other one cannot exceed 50% of the total.
  - If we allow this configuration, it precludes having a fixed allocation with more than 50% to any interface.

Possible paths forward

A. **Split the error budget evenly between C2M and C2C**
    - If COM will be used for both, then DER_0=1.33e-5 for both
    - Margin left on the table in many use cases

B. **Allocate the whole error budget to the C2M**
    - If COM will be used, then DER_0=2.67e-5
    - C2C can still be used in an xGMII Extender with a larger error budget
    - Not friendly for copper cable

C. **Conditional allocation – different C2M BER spec based on whether C2C is used or not, and vice versa**
    - Two sets of C2M specs (with different BER) for both hosts and modules

Comparison of conditional and fixed allocations