All-lanes error counting

(comments R2-198 and others)

Piers Dawe
NVIDIA

Comment R2-198

- R2-198 SC 180.9.14 P 492 L 16 TR
- 802.3 specifies equipments by their behaviour at external interfaces. It does not require a particular implementation; all implementations should be held to the same standard so that they are interoperable and the spec does not exhibit favouritism. Additionally it provides example specifications for some internal interfaces such as AUI C2M, but these are not required.
- D3.2 erroneously disallows methods 174A.9.7 (BLER method for all lanes using PMA measurements) and 174A.11 (using PCS measurements). The former is incorrect because errors are shared by a common FEC), and also, measuring an individual lane and extrapolating as far as 6.9e-28 is subject to errors that can be reduced by pooling the lane measurements as described in 174A.9.7. The latter is incorrect because it's at the PCS where the errors matter, and there may be no AUI in a product.
  Compare 174A.10, Error ratio tests for 800GBASE-LR1... where "the number of lanes p = 1" although there are four signals rx_signal_xi, rx_signal_xq, rx_signal_yi, rx_signal_yq whose errors could have been counted separately. That might be interesting for diagnostics but it is not appropriate as a spec limit.

  In general (for the receive side), we should consider errors at or near the PCS first, and fall back to counters nearer to the PMD when necessary. Because this is an equipment spec not a component spec.

  As explained before, receiver implementers who don't like the all-lane methods can conservatively use 174A.9.5 or 174A.9.6, if they allow for extra scatter. Qualifying a module for 8 x DR1 is a different matter anyway (different BLER block time on a lane).
  The all-lanes calculation is still conservative because each receiver lane is tested with minimum victim and maximum optical crosstalk aggressor, and they can't all be like that, and because each transmitter lane is specified separately so they are all at least as good as minimum that the receiver is tested against.
- Revert the changes that eliminated 174A.9.7 and 174A.11 in the IMDD clauses.

See also

- Comment R2-191 on BLER extrapolation and rubin_3dj_01_2609
- Comment R2-143, a valid 800GBASE-R signal is suitable
- R1-274, R1-275, R1-277, R1-294, I-419, I-420, I-421 and I-493

Errors are shared across lanes by the FEC

- 800GBASE-LR1 allows 1.5 dB power difference between X and Y polarizations and 1 dB I-Q amplitude imbalance
- So one sub-lane can have many more errors than another
- The most errored sub-lane might be 3x the average
  - That's OK
- Just like in IMDD, there is no limit on the distribution of errors between LSB and MSB
  - That's OK too

Example

- A receiver with 4 lanes, with a standard distribution of sensitivity of 0.1 dB
  - Some lanes will have 50+% more raw errors than others
  - These lanes will dominate post-FEC performance
  - But the lack of errors from the better lanes is real and should not be ignored
- Per-lane specification wastes a substantial proportion of the lane-to-lane standard deviation