HOW TO ACHIEVE 3 M NO FEC

Adee Ran, Intel Corp.

Supporters:

Rich Mellitz, Intel Corp.

Omer Sella, Mellanox

Oded Wertheim, Mellanox

Background

- During the last couple of meeting we have been discussing the desire to include “3 m no FEC” required performance in 802.3by.
- A market motivation for no-FEC operation was demonstrated.
- Following straw poll in July meeting, work was done to explore possible ways of addressing this desire. Most of it involved changing TX and/or RX specs.
- To change or not to change?
  - Activity since last meeting doesn’t seem to create a strong consensus to change the specs in a specific way.
  - Based on the proposals discussed in the ad hoc, can we address this question first, before going into the details?

Is there a solid case?

- The whole discussion is based on the premise “We saw it work ➔ there is margin”
  - Although… detailed compliance data of the channels and parts that were used to make this claim was not presented to date
  - Claims of “hidden margin” in the COM method (or tool) were refuted.
- A reasonable explanation is that the parts that were used were <u>better than minimally compliant</u>
  - Does that mean that there is extra margin to be harvested?

Statistical vs. worst-case argument

- Our specs are written to ensure adequate performance of a system where TX, cable, and RX are all worst case compliant
  - Using the word “shall”
  - Vendors typically take these very seriously
- The combination of worst-case everything is statistically very rare. Usually one or more of the parts is much better than worst case.
- Manufacturers typically screen parts in attempt to prevent bad (possibly incompliant) parts from reaching the market.

Not all parts are created equal

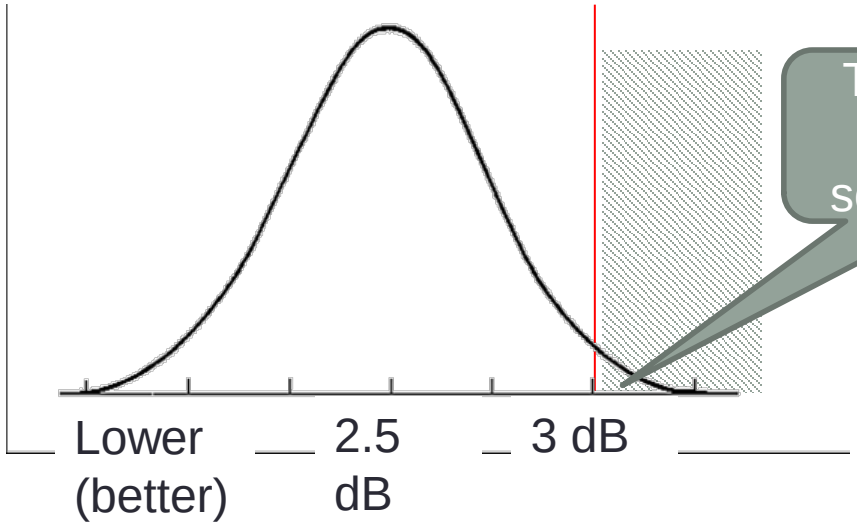

Receiver “channel tolerance” (in interference tolerance test)

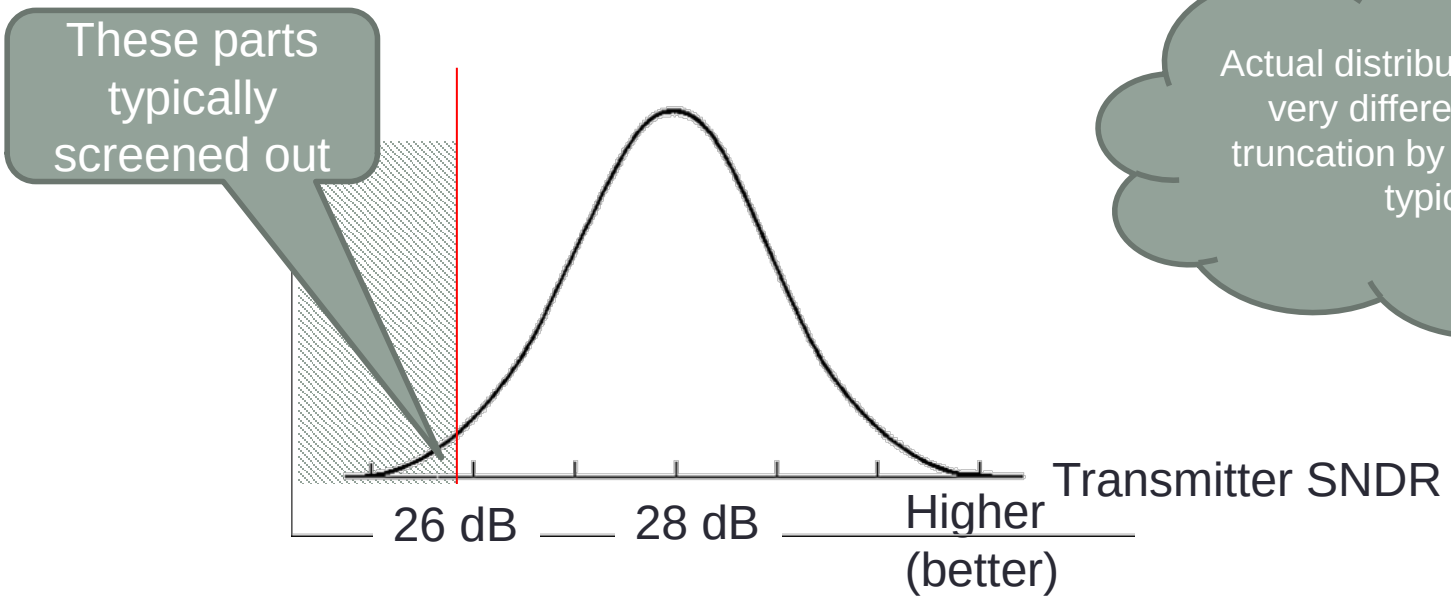

Actual distributions may be very different, but the truncation by screening is typical.

(Similar distributions for jitter, amplitude, etc.)

Consequences of statistics

- The result is that practical systems are expected to have margin.
  - This is a good thing!
- Some people rely on that and demand more
  - “We want BER<1e-15, not just 1e-12”
  - “We need 140 meters of Cat 5e, not just 100”
  - And in this case “we need 3 meters”
- But if these were the specs, a larger portion of devices would have to be screened out.

Possible consequences of changing specs… lower yield

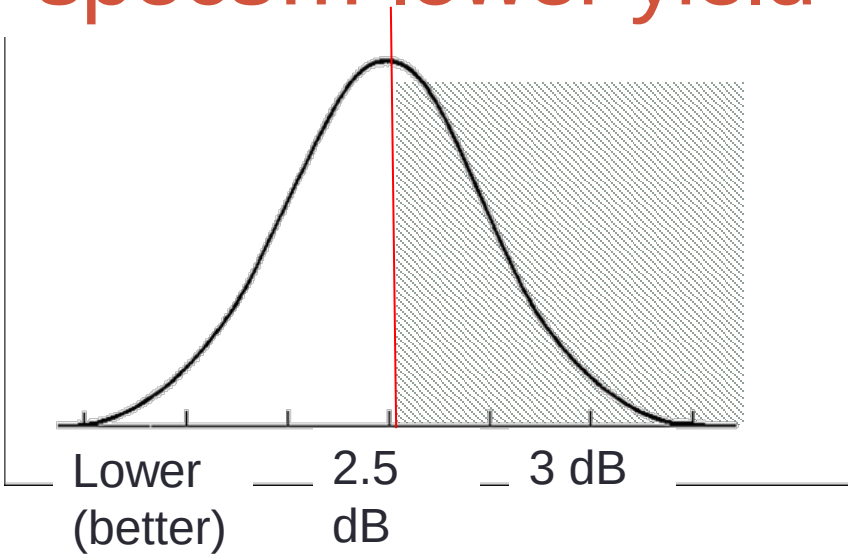

Receiver “channel tolerance” (in interference tolerance test)

Yield has implications on cost

Real distributions and probabilities may be different.

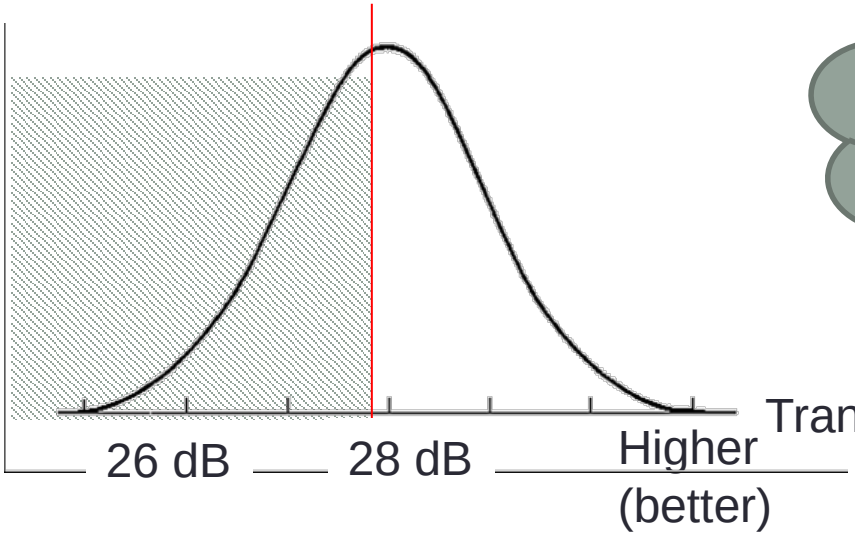

Transmitter SNDR

… or redesign SERDES for stricter specs?

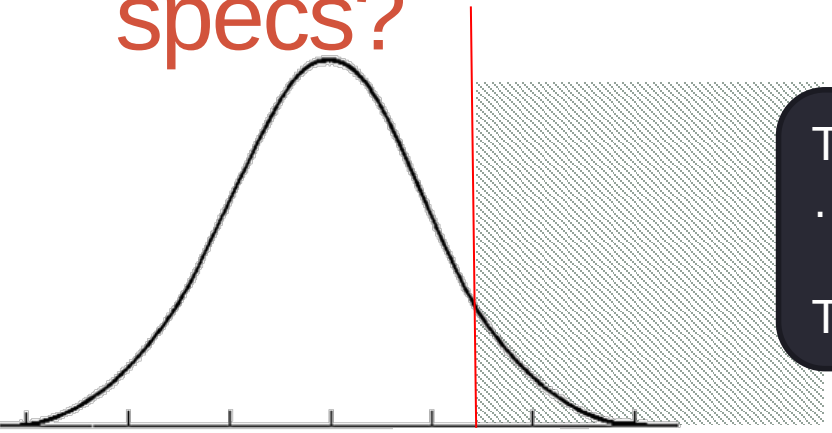

Lower (better)
2.5 dB
3 dB

This is not what this task force started with!
… also likely to have power and cost implications

THIS IS UNLIKELY

Receiver “channel tolerance” (in interference tolerance test)

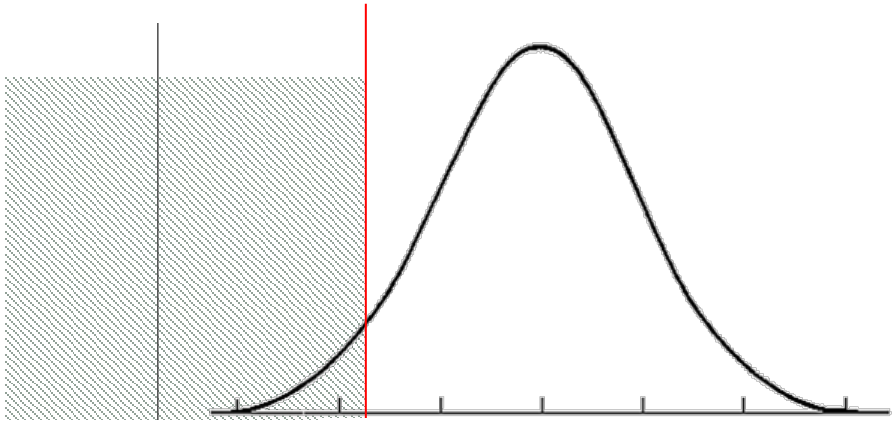

26 dB
28 dB
Higher (better)
Transmitter SNDR

“BUT IT WORKS”

Does it actually work as we expect?

- EA Plugfest report (dambrosia_3by_01_0915) suggests that only 90% of the no-FEC tests combinations achieved the specified BER…
  - Even lower success levels on cables with IL>20 dB
  - The parts used were likely typical or better
  - Temperature, humidity etc. were not stressed
- It seems to suggest that “3 m no FEC” doesn’t actually work with a wide variety of receivers and transmitters, while 2 m seems to work more robustly – which agrees with the current COM parameters and budget allocation.
- Internal stressed measurements done in Intel show similar results, and support this conclusion.

Layers of challenge

- Following an observation made during the last ad hoc call… there are several layers of challenge in building a product:

1. Show an analysis or simulation that a system (TX+channel+RX) works as expected
2. Build a lab demo that shows desired performance is met (e.g. BER on a link).
3. Pass all compliance tests with a specific part.
4. Pass compliance tests with N parts to build confidence that the design is robust.
5. Have a screening process that ensures most of the parts you ship would pass compliance tests (without throwing away too many good ones)

NOW WHAT?

So how can 3 m no FEC be achieved?

What options exist for 3 m no FEC – without changing the specs?

- From a PHY vendor point of view – either/or
  - Screen parts and choose high-performance ones (that tolerate 2 dB COM, etc.) for customers that demand it
  - Design a better RX (e.g. LFEQ, better DFE, etc.) – as in proposals – and have a competitive advantage
- From a system vendor point of view – either/or
  - Choose parts which have better than minimum performance
    - Lower TX jitter, higher SNR
    - Better RX tolerance (smaller link budget consumed)
    - That has been widely done for BER<1e-15
  - Use better host PCBs (lower loss, impedance matching)
  - Use better (thicker) cables
  - Rely on statistics.