Resolving COM and Rx ITT concerns

Adee Ran

Intel Corp.

Background

- COM procedure for testing channels uses a specific set of package model and device termination values.
- The receiver interference tolerance test (RITT) is calibrated using a procedure based on COM.
- A receiver under test may have package and termination different than the model used for COM. Likewise for the transmitter in the test setup.
- We have been discussing possible scenarios where combination of borderline receiver and channel to not operate as expected.

Goal of this presentation

- State the problem scenarios.
- Discuss implication in real life.
- Consider possible options.
- Recommend a path.

Problem statement

- A channel C (e.g. cable assembly) barely passes COM (assuming a certain termination impedance) and is deemed compliant
- A receiver R has a different termination, and is tested with a “bad” test channel that requires little noise as calibration; the receiver is actually matched to the test channel so it is tested with less stress than other channels would cause, and barely passes RITT, so is deemed compliant
- The combination of channel C and receiver R is mismatched, and fails to interoperate.

This was nicely visualized in hidaka_060717_3cd_adhoc…

Hyperspaces of channels…

Test channel strictly worse than all compliant channels

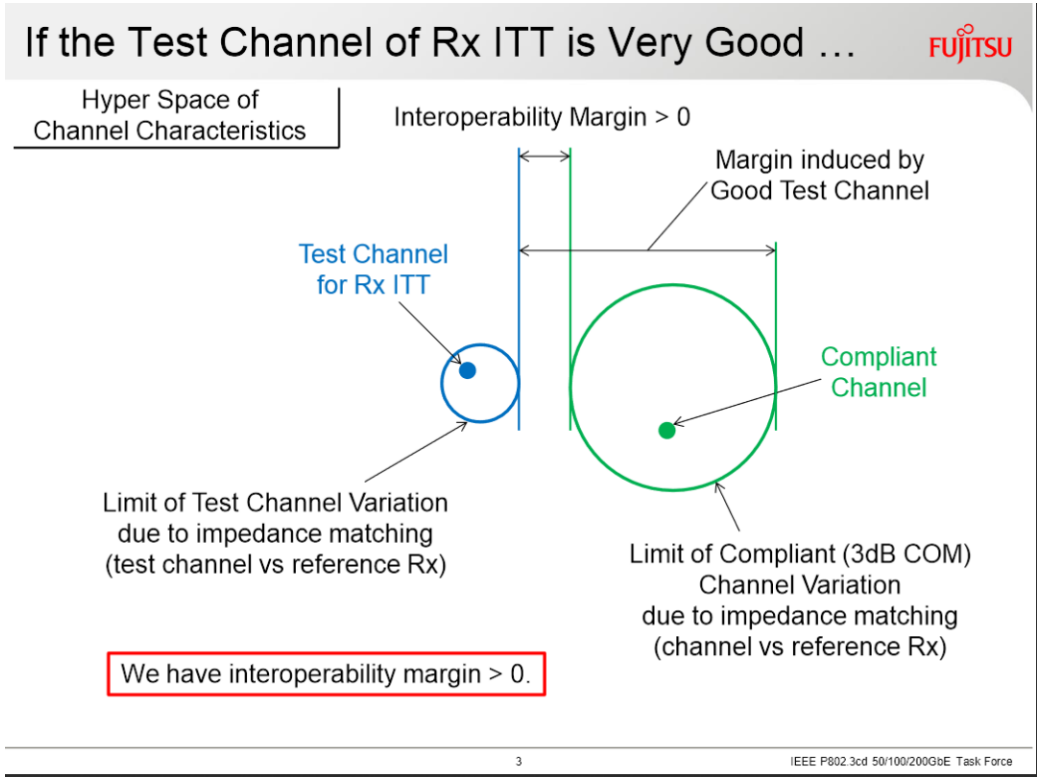

Test channel somewhere in the space of compliant channels

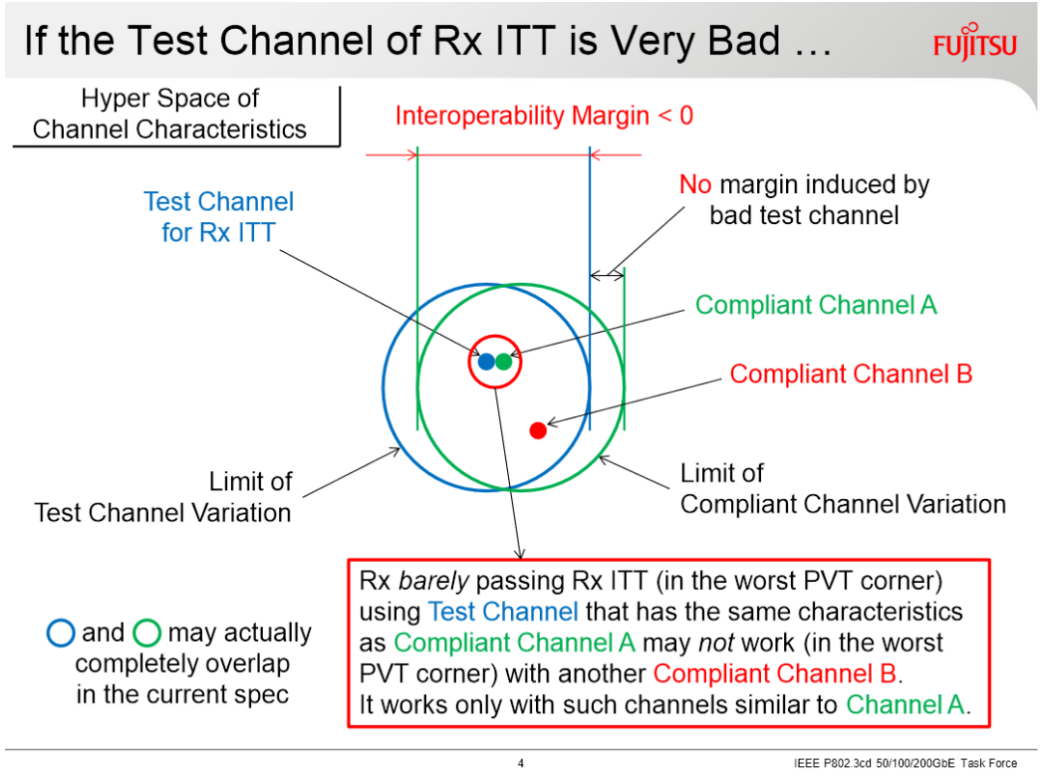

Continuing this metaphor… (1)

When we test receivers, the RITT channel must be taken from this complementary subset

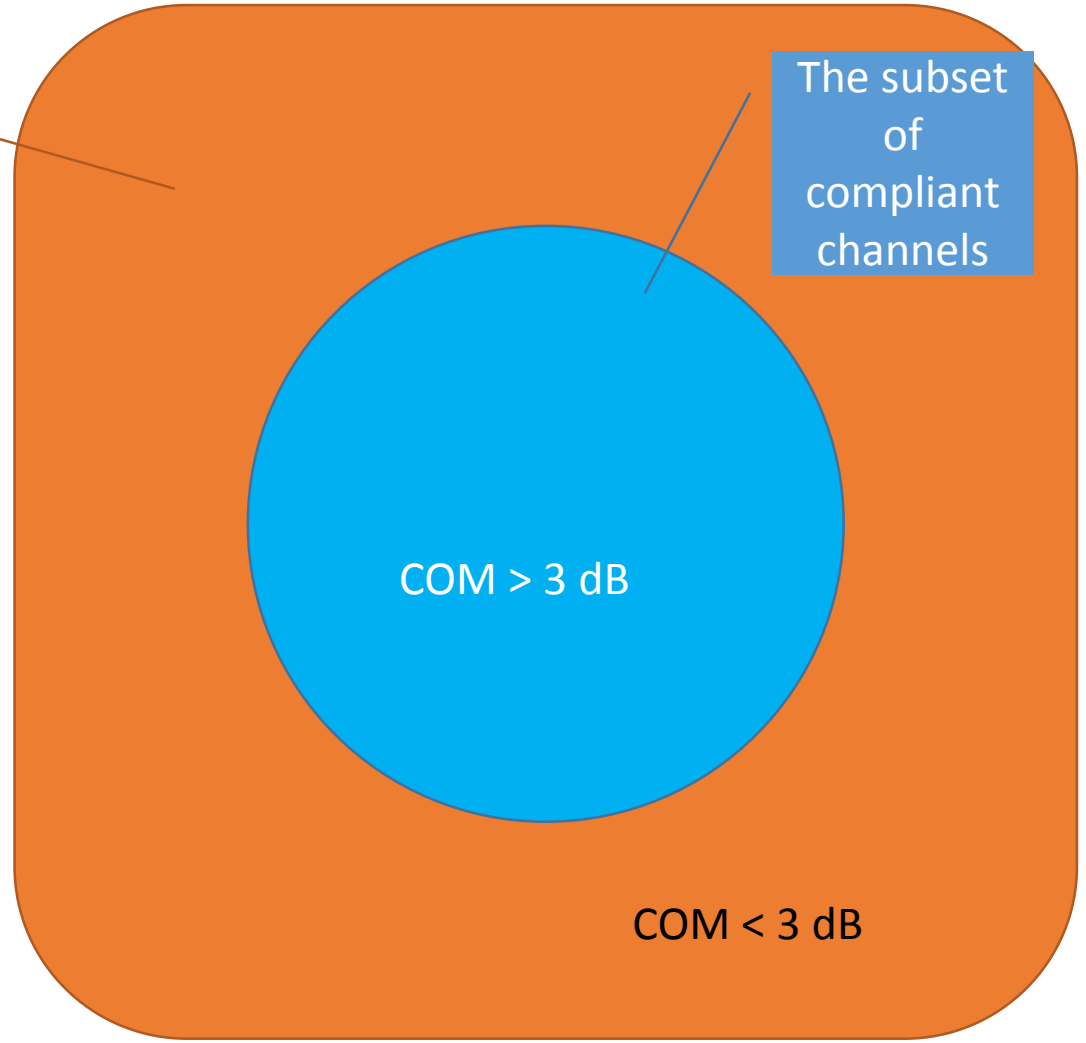

Continuing this metaphor… (2)

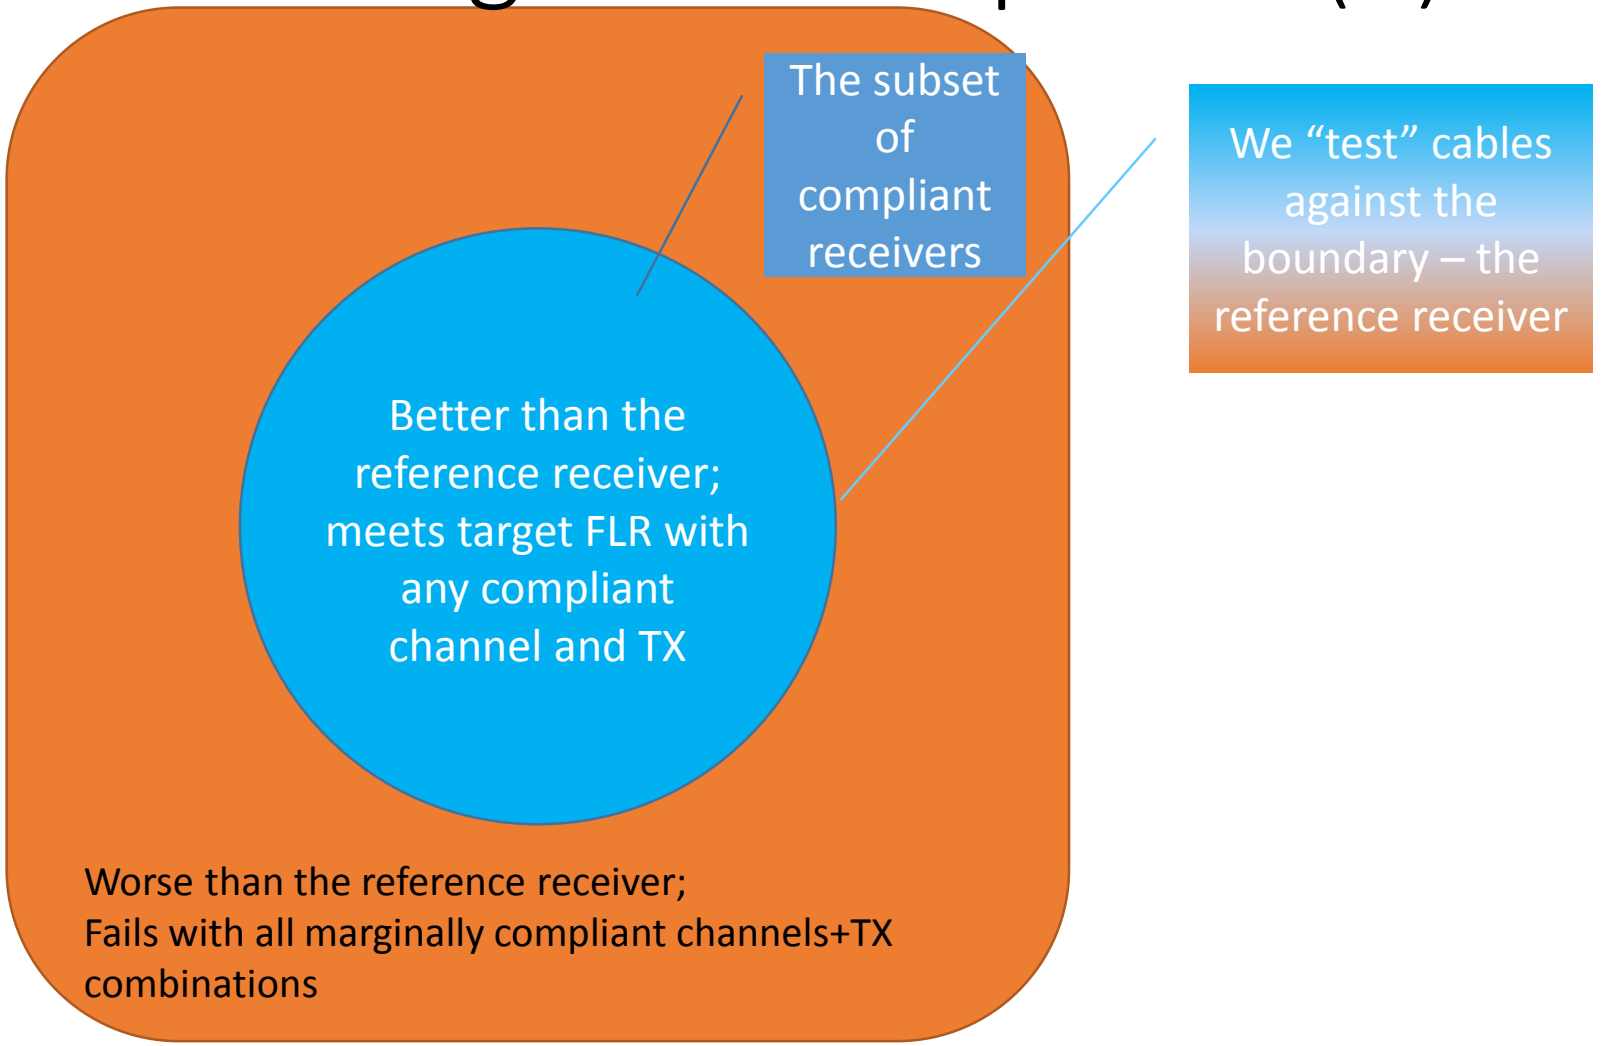

Continuing this metaphor… (3)

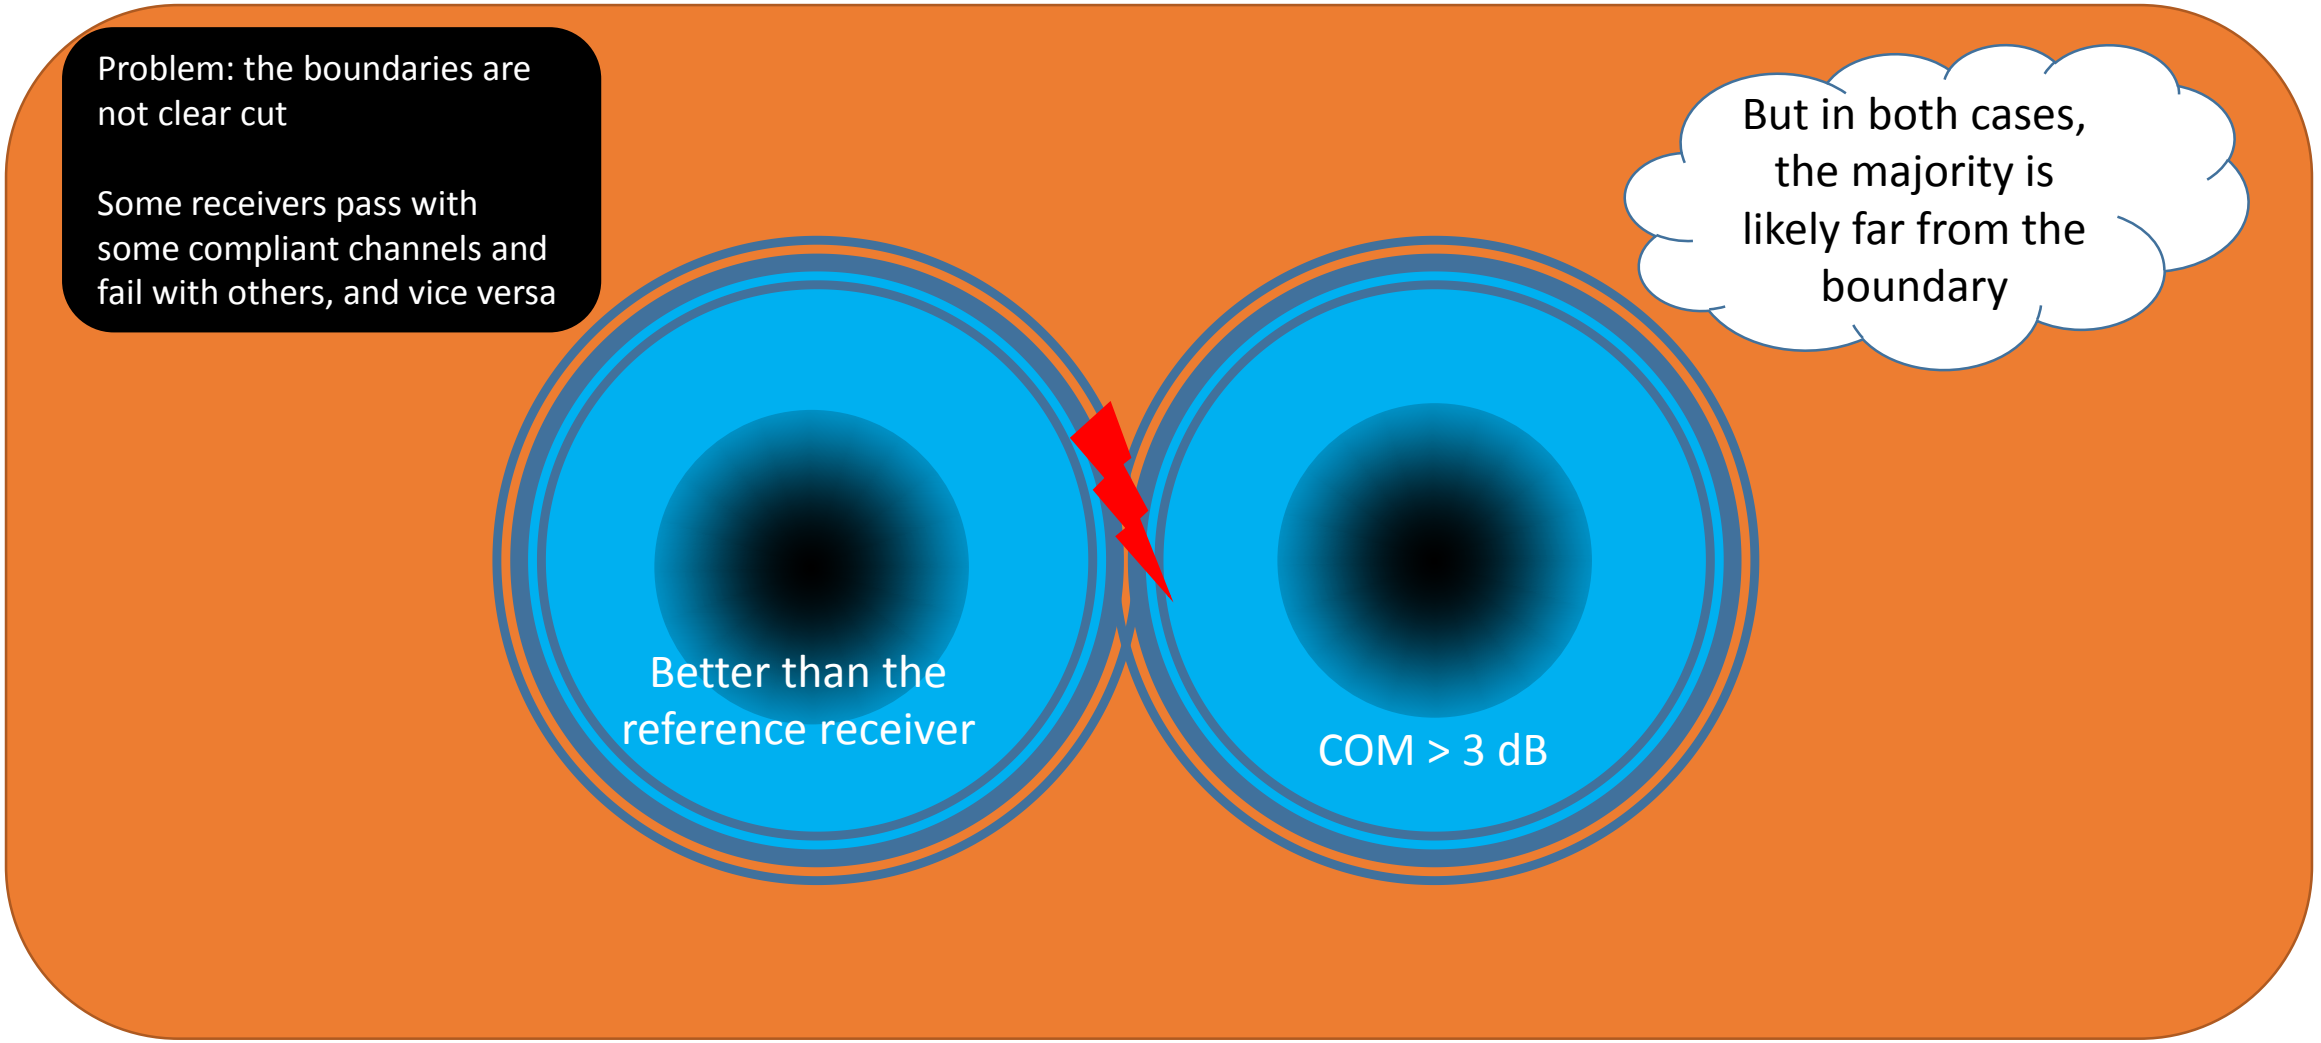

COM is no prophecy

- The fact that COM is not a perfect predictor is not new.
  - Receiver performance may depend on many fine details that we don't account for…
    - The exact jitter frequency profile and PDF
    - The positions of the reflections in time
    - The exact values of the ISI that the (quantized) equalizers have to mitigate
    - Design details.
  - COM also does not account for all possible forms of jitter and nonlinearity, and uses some approximations and arbitrary choices of parameters.
- Accounting for more phenomena and combinations would unnecessarily encumber the standard, would burden the designs and would make the compliance tests even more complex.
- COM is not a signal integrity simulator and does not predict actual performance.

There is no inherent margin

- The fact that the standard does not include an interoperability margin is not new.
- Even if COM, Tx specs and RITT were all perfect predictors, there could theoretically be receivers, transmitters and channels that pass their compliance tests and when plugged together operate at just the normative performance level (say a frame loss ratio of 6.2e-10) but not any better.
- Would that be okay?

Margin is vendor's call

- If vendor X sells parts that fail in some customer's system, while vendor Y's competing parts work well, then…
  - X might lose a deal to Y
  - X might get the deal but get angry phone calls from users (or worse).
- X can't test each part with each possible combination of other parts
- X will likely design with some internal margin to avoid customer dissatisfaction, have good yield, and be competitive
  - In fact, some customers demand performance better than the standard requirements (e.g. BER much lower than 1e-12 is a common requirement).
- More margin is good… but can add cost, power, schedule
- Less margin is a risk… but risks are sometimes taken

HMME?

- How Much Margin is Enough? That is the question:
  - Receiver should be better than the reference receiver…
  - Channels should be better than 3 dB COM…
  - Transmitters should be better than worst case in all specifications…
- … By how much?
- It is for vendors (and sometimes customers) to decide, so that expected performance is achieved at reasonable cost and risk.
- It is not our job as a standards body to dictate the answer.
  - In fact the details of the problem are so complex that many people won't understand what margin there is; any margin we dictate may become the baseline that the market will try to exceed.
  - This would effectively add unnecessary cost without any risk reduction.

What are our options?

1. Add more test cases to make the boundary more clear cut.
   - The additional coverage is questionable. We can't guarantee interoperability.
   - This will add burden to validation and testing.
2. Add explicit margin, e.g. by using more stress in RITT or using worse parameters in COM.
   - This happens anyway, it's not our job, and it will add cost
3. Adopt new methods of characterization, e.g. TDR-based specs.
   - That can always happen, but we should use the best method at hand
4. Continue discussions…
   - We should converge
5. Decide to move forward
   - How?

Moving forward - proposal

- Use the existing COM methods (as in D2.0)
- Use the current parameter values and start examining example channels
  - Solicit more channel contributions
  - Maybe we should change some parameters
- Solicit test results of key parameters that we defined
  - Transmitter SNDR, SNR_{ISI}, pulse peak to steady state ratio, jitter
  - Receiver ITT/JTT results with calibrated setups
  - BER (or SER, or more detailed) system performance
- This kind of data was presented in 802.3by and was highly effective

Thank you

Questions?