TDECQ, slowness, badness and overshoot (revised)

Piers Dawe

Mellanox Technologies

Supporters

• Marco Mazzini	Cisco
• Johan Jacob Mohr	Mellanox
• Santiago Echeverri-Chacón	Mellanox

- *New slides: 8, 14, 11, 15, 18, 21, 22*

Introduction

- In the TDECQ method, there are several specs on a signal
  - Four limits in the in-force 802.3cd and 802.3cn
- We have four limits to protect the different parts of the receiver, and the link quality, against different threats
- Another limit, against overshoot, is in IEEE P802.3cu draft 2.0, but one of the four has been partly removed
- This presentation addresses their different characteristics and uses, showing that all four, and probably overshoot too, are separate and needed
- Also presents a way of measuring overshoot

Cost of correcting + what's not corrected = penalty

Slowness + badness = penalty

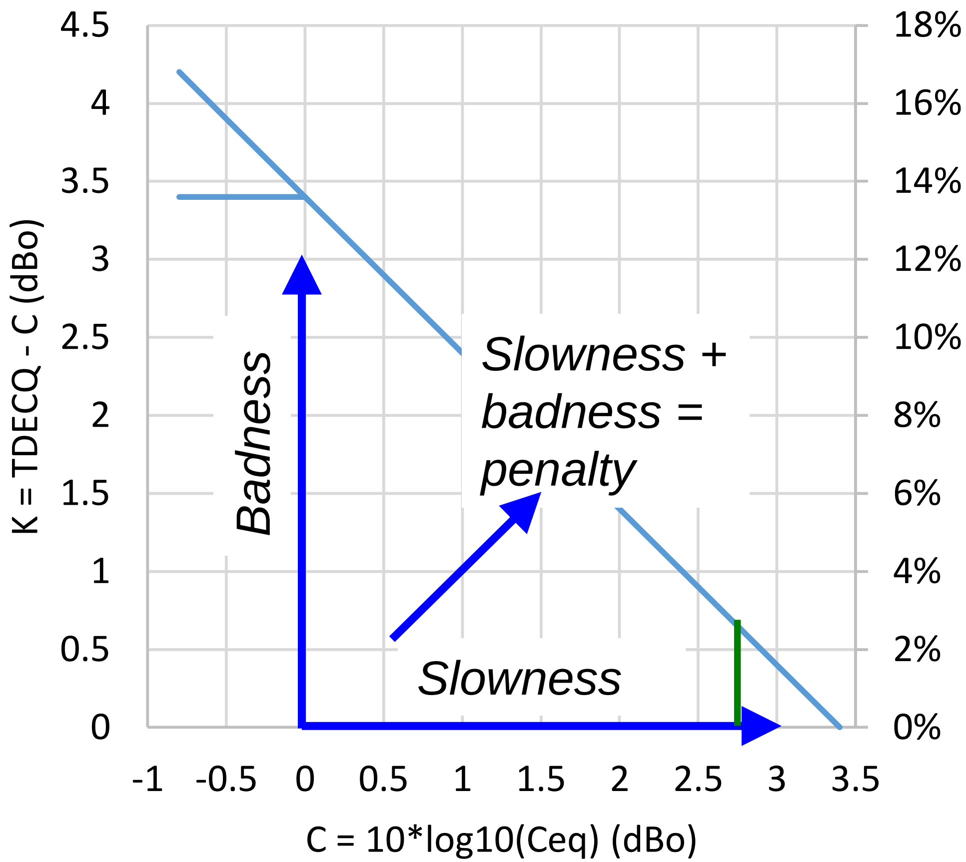

- TDECQ map was introduced in [dawe_3cd_01a_0318] (nearly two years ago)
- We have used different names over time, e.g.
- Slowness penalty, noise enhancement [penalty] or **C** for "slowness":
- **C is the cost of correcting what the reference equalizer can correct**
- roughness and signal's noise penalties, residual ISI and noise penalties, residual eye penalty, or **K** for "badness";
- **K is what isn't corrected**

See:

Transmitter and Dispersion Eye Closure Quaternary (TDECQ) and Its Sensitivity to Impairments in PAM4 Waveforms
Santiago Echeverri-Chacón, Johan Jacob Mohr, Juan José Vegas Olmos, Piers Dawe, Bjarke Vad Pedersen, Thorkild Franck, and Steen Bak Christensen, Journal of Lightwave Technology, **37**(3), pp 852-860, February 1, 2019

Examples (as before equalizer)

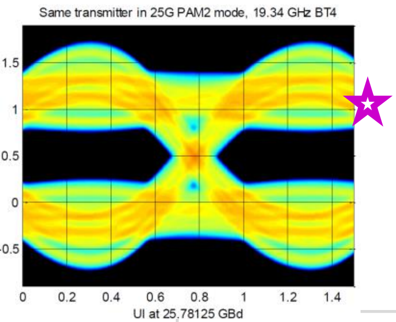

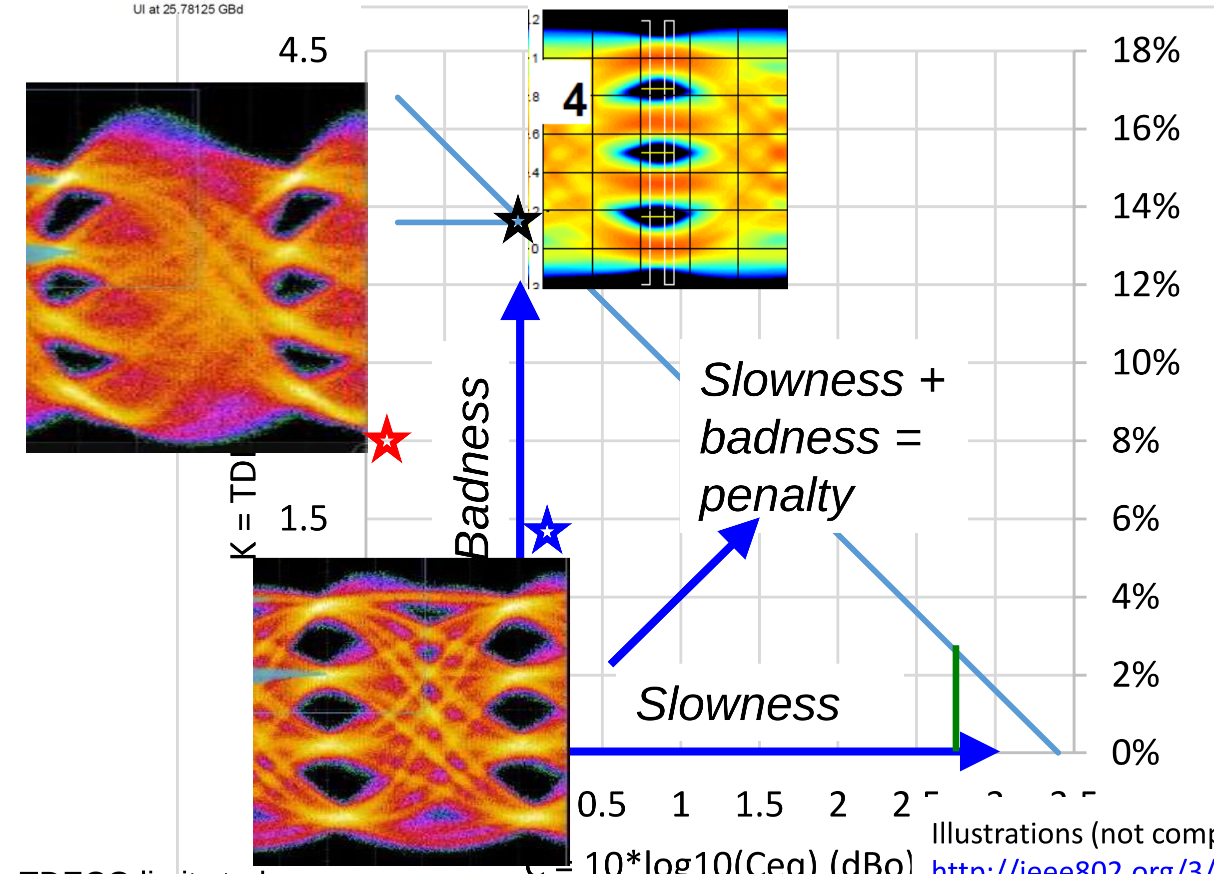

- Four limits at left, top, diagonal-right and right
- Too much: emphasis, bad, penalty, slow, respectively

Alternative limits, same measurement

- This presentation concentrates on two of the limits, as shown in blue

Illustrations (not comparable with each other) from
http://ieee802.org/3/cd/public/July18/dawe_3cd_01b_0718.pdf#page=11
http://ieee802.org/3/cd/public/July18/dawe_3cd_01b_0718.pdf#page=9
http://ieee802.org/3/cu/public/Jan20/cole_3cu_01b_0120.pdf#page=14
http://ieee802.org/3/cu/public/Jan20/cole_3cu_01b_0120.pdf#page=12

TDECQ limits today:
100GBASE-DR, -FR1 and -LR1 and 400GBASE-FR4: 3.4 dB
400GBASE-LR4-6: 3.5 dB
K limit for 100GBASE-DR is 3.4 dB
K limit for others is debated

Compare a non-equalized PMD

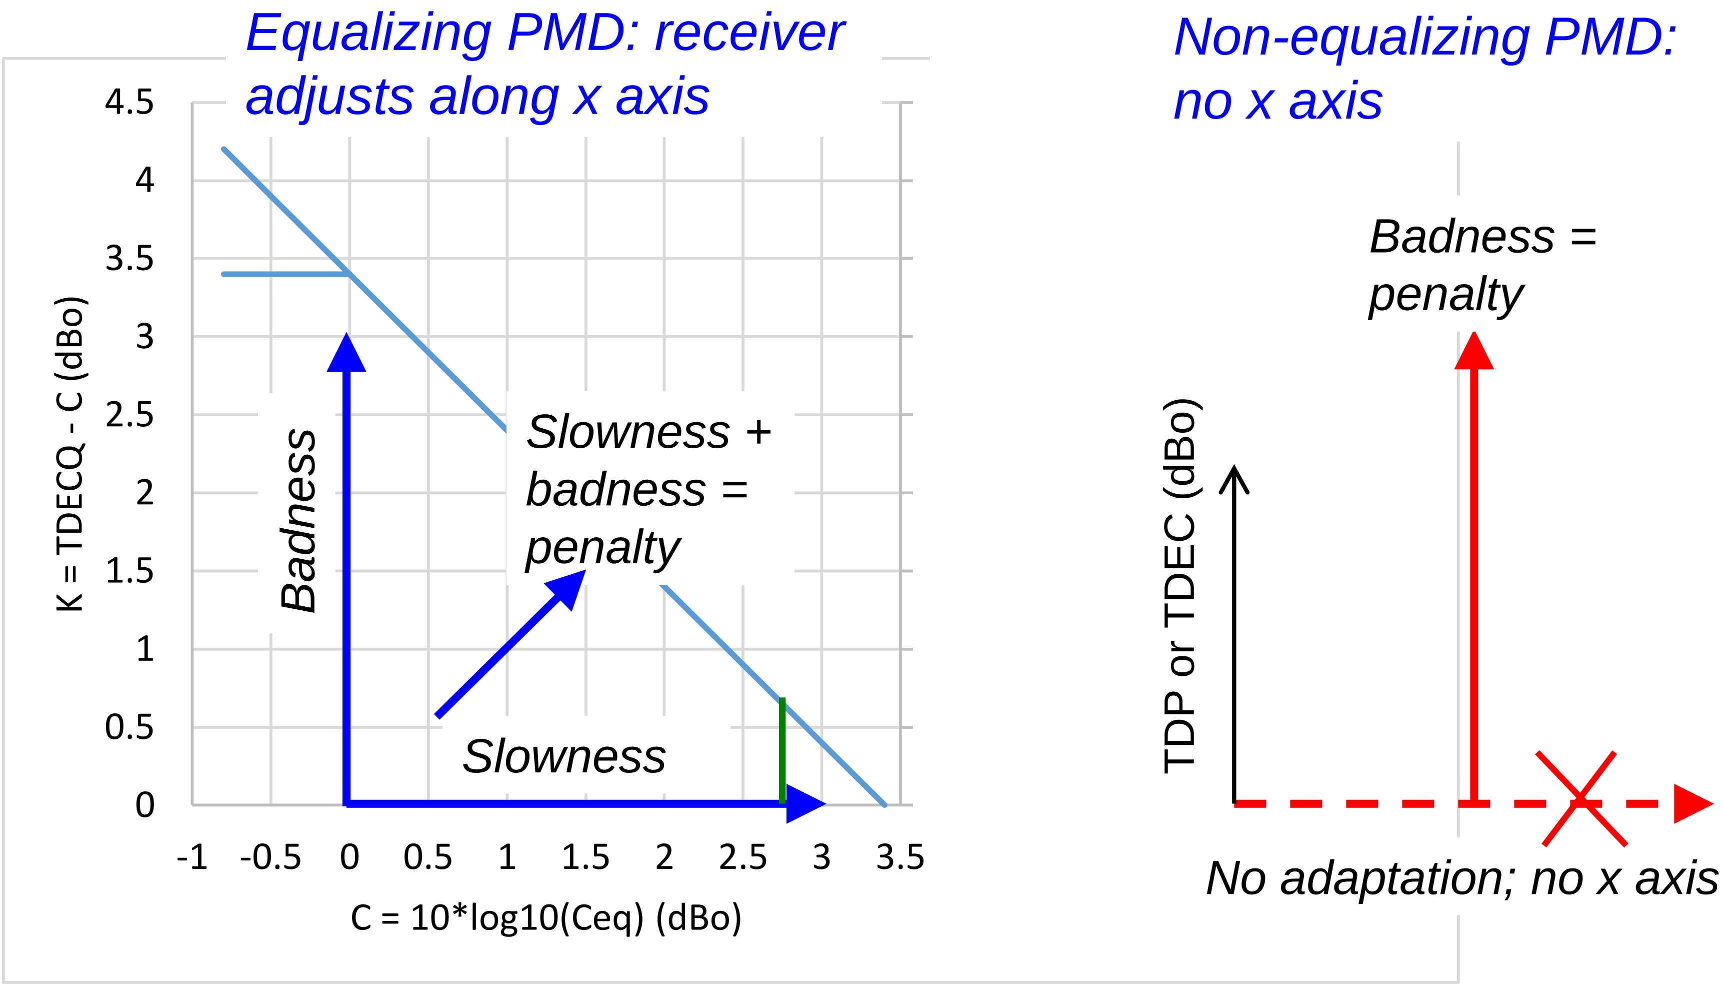

We always limit badness

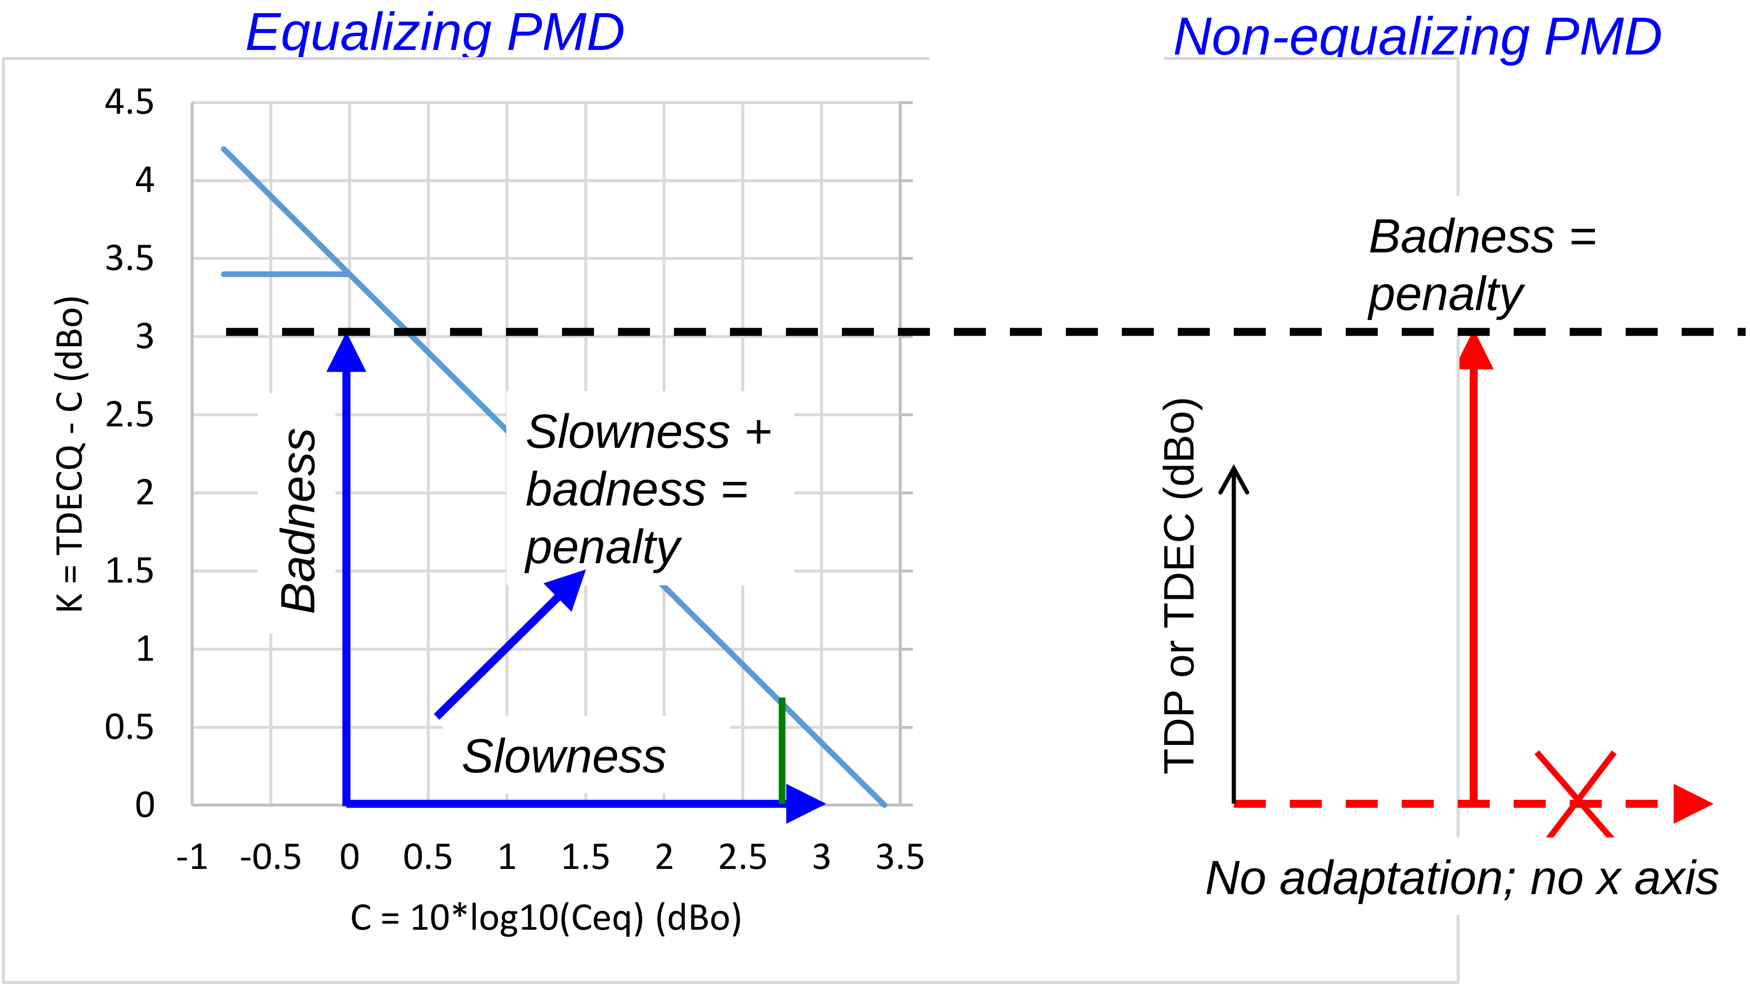

ITU-T limits badness too

Equalizing PAM4 PMD

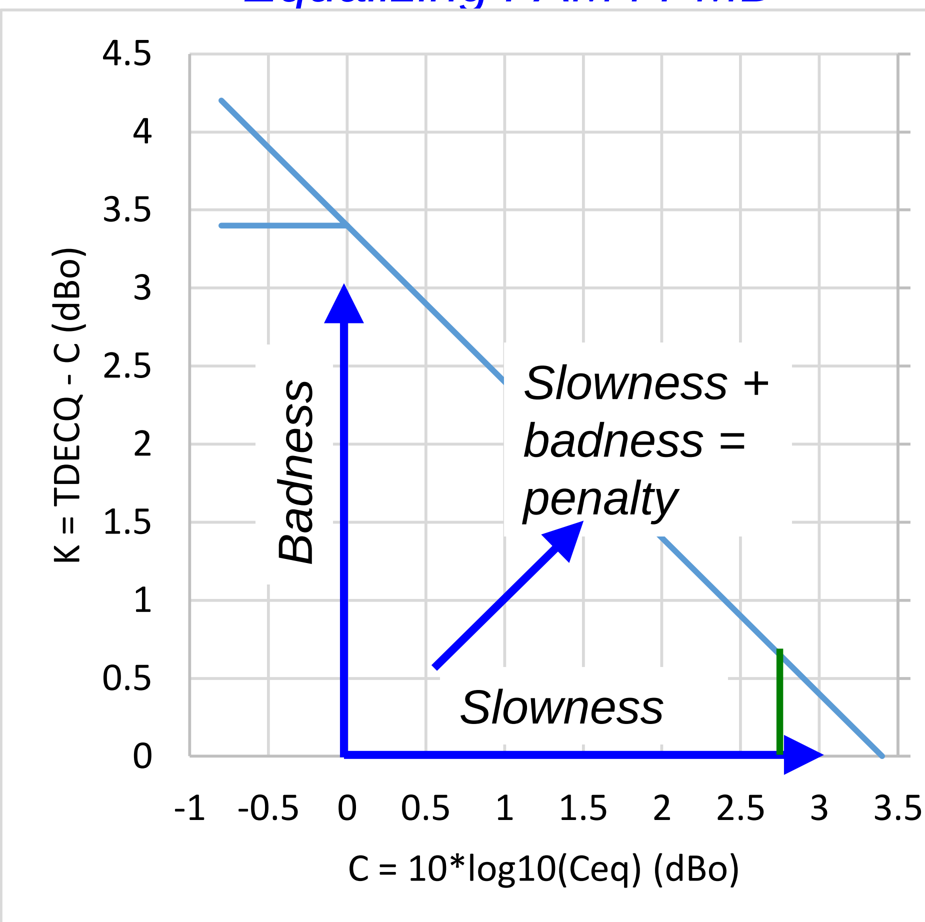

Equalizing DP-DQPSK PMD

EVM (error vector magnitude) as defined in Recommendation ITU-T G.698.2, is also a measure of badness ***after the equalizer***

All in-force optical PAM4 is done this way

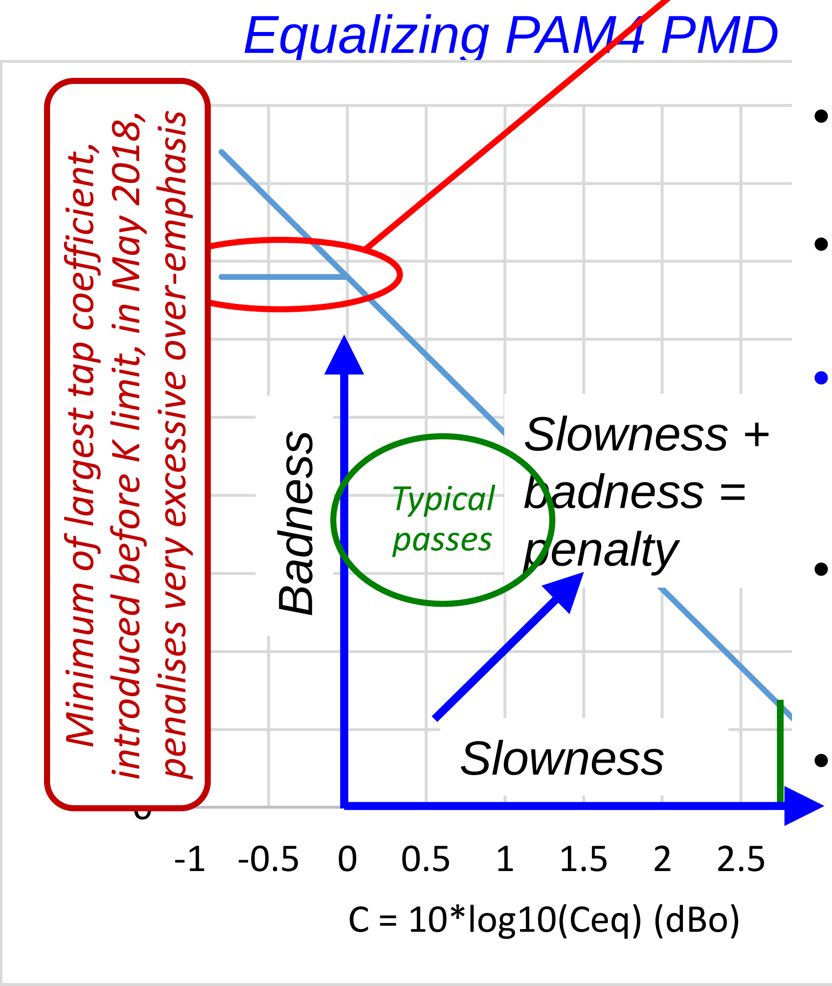

- Signal can have +ve or –ve C
  - +ve C (slower than neutral) is usual
- Budget is calculated correctly for either sign of C
- K is limited to the same limit as TDECQ
  - which depends on PMD type
- **K protects receiver and helps avoid bad error floor**
  - **See slide 11**
- Protects against non-equalizable elements of signal
  - including non-equalizable overshoot
- K is not intended to protect against over-emphasis
  - Another spec protects against equalizable over-emphasis: see dawe_3cd_01b_0518
- TDECQ, C and K from one measurement: **no Tx measurement cost impact**

All in-force optical PAM4 is done this way

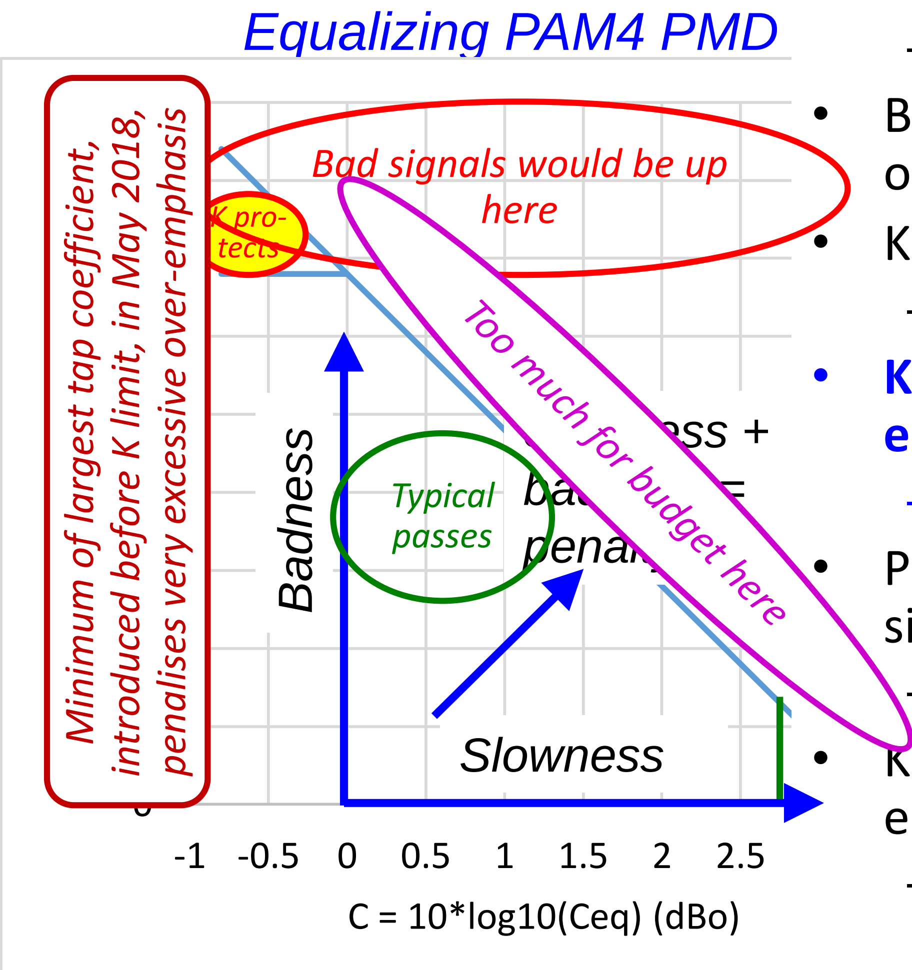

- Signal can have +ve or –ve C
  - +ve C (slower than neutral) is usual
- Budget is calculated correctly for either sign of C
- K is limited to the same limit as TDECQ
  - which depends on PMD type
- **K protects receiver and helps avoid bad error floor**
  - **See slide 11**
- Protects against non-equalizable elements of signal
  - including non-equalizable overshoot
- K is not intended to protect against over-emphasis
  - Another spec protects against equalizable over-emphasis: see dawe_3cd_01b_0518
- TDECQ, C and K from one measurement: **no Tx measurement cost impact**

Allowing unreasonable signals adds a burden to receiver qualification

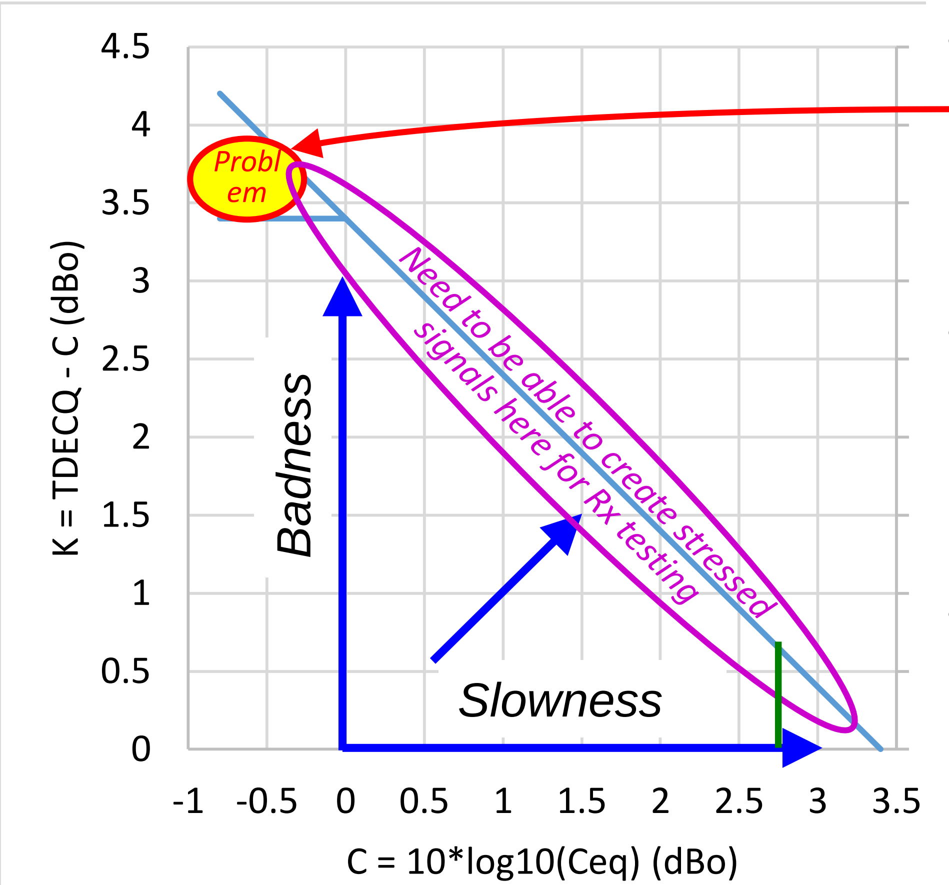

- If we allow signals in the top left corner we have to make "stressed eyes" so we can test them
- "Very difficult region to verify SRS compliance":
  - http://ieee802.org/3/cd/public/July18/mazzini_3cd_01d_0718.pdf
- So there is a measurement cost impact; removing the K limit **would add** to receiver measurement costs

At least two sorts of overshoot

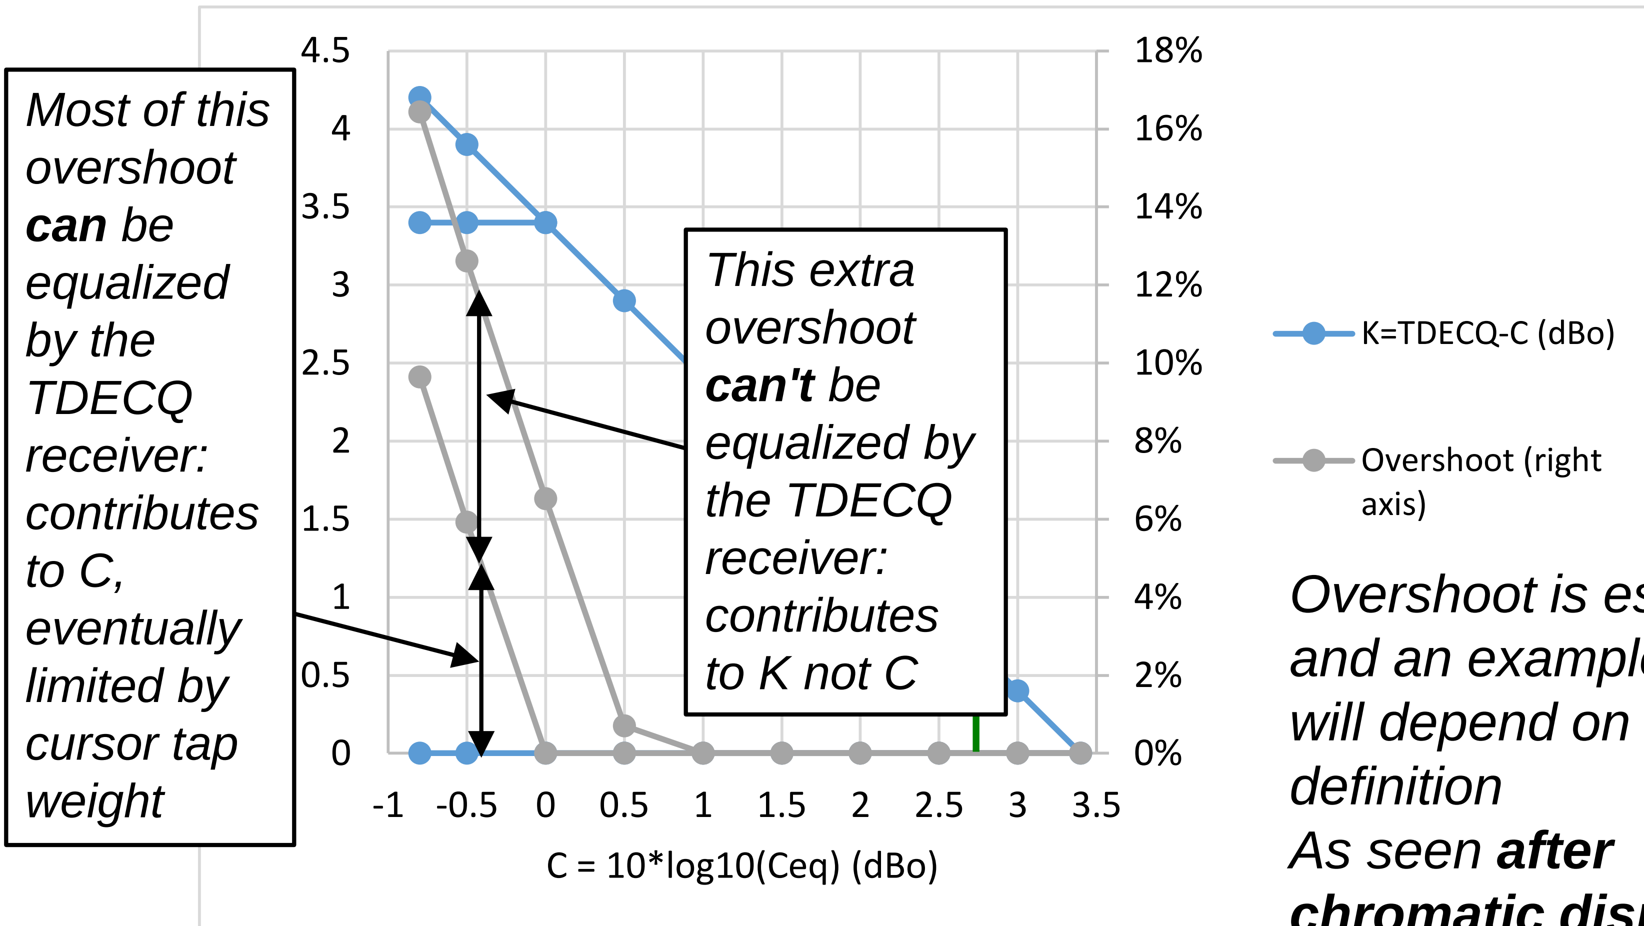

Overshoot is estimated and an example only; will depend on detail of definition

*As seen **after chromatic dispersion** by TDECQ receiver*

Bad signals with equalizable overshoot were studied in 2018

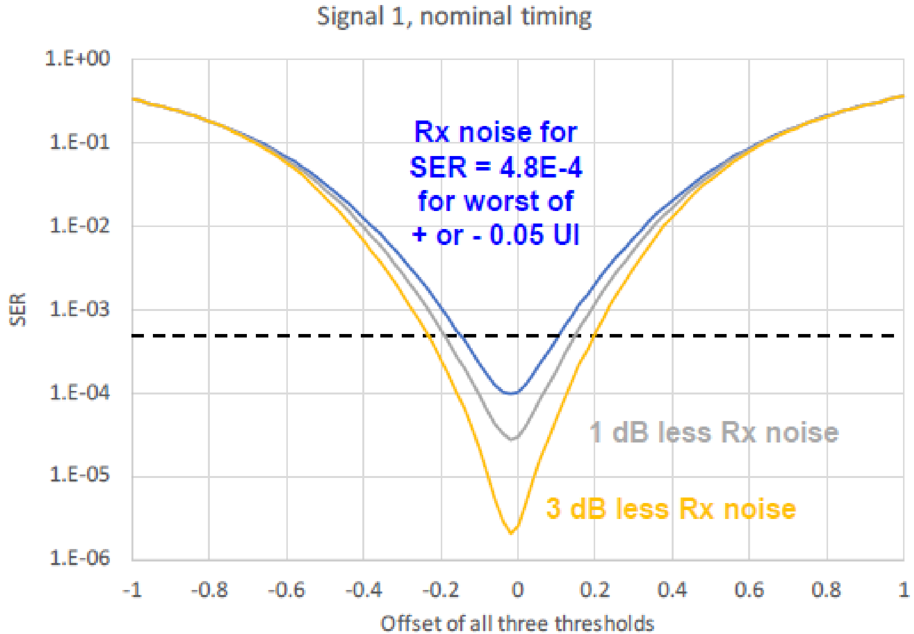

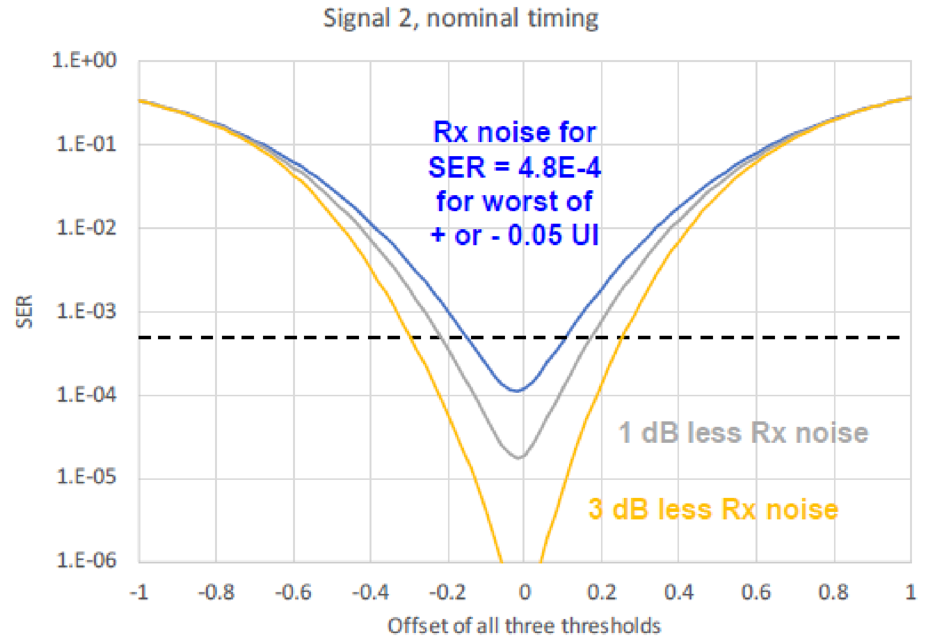

Requires better threshold setting, error floor is worse

Better because K is lower ~ 3 dB

- Both signals have 3 dB TDECQ
- Signal 1 is over-emphasised, signal 2 is near neutral
- This is from http://ieee802.org/3/cd/public/adhoc/archive/anslow_062718_3cd_adhoc.pdf
- Signal 1 has worse error floor; real-world impairments will exacerbate the difference
- So would a higher TDECQ limit
- http://www.ieee802.org/3/cd/public/July18/dawe_3cd_01b_0718.pdf slides 9 and 10 show the same effect

Bad optical signal and error floor

K should limited to 3.5 dB (e.g. in 400GBASE-LR4-6), like TDECQ
Without K or overshoot limit, it could be up to 4.3 dB

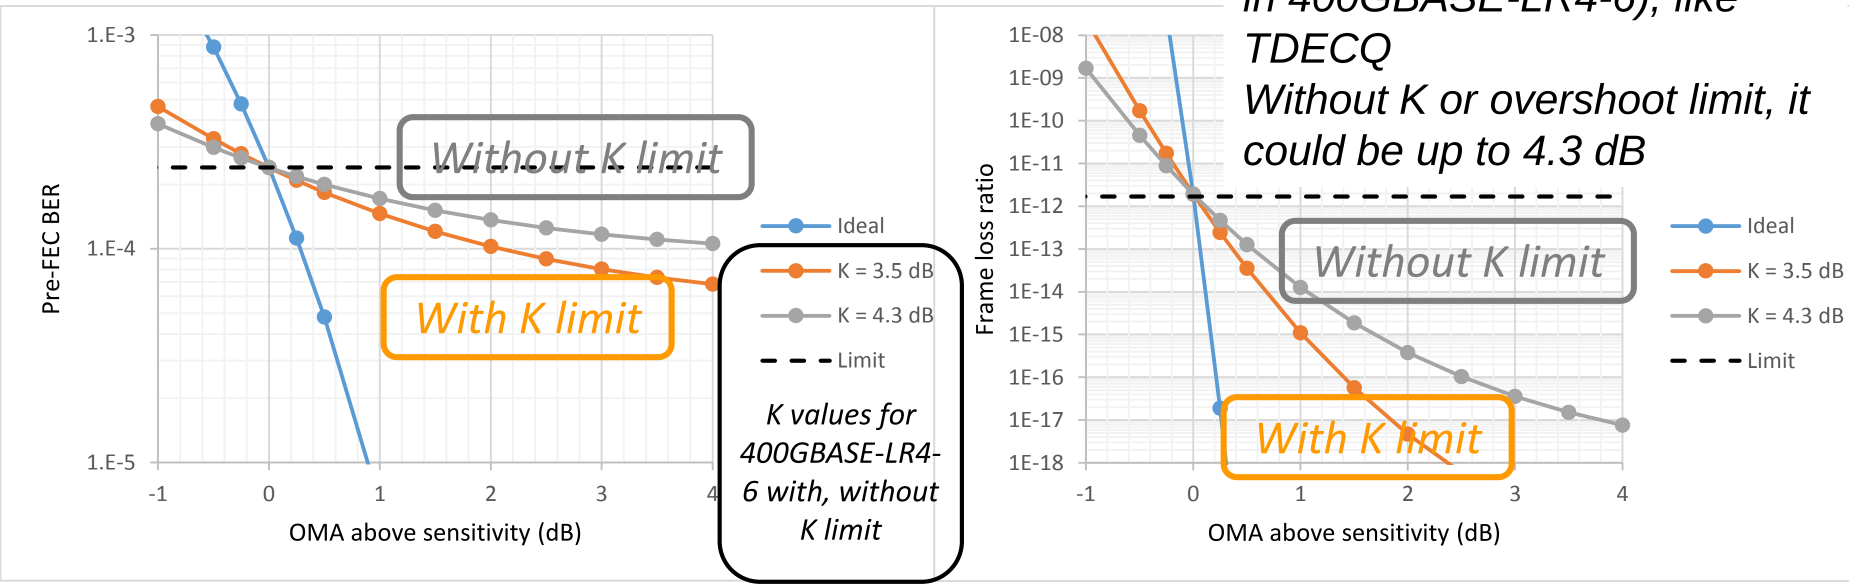

- Textbook theory, assuming Gaussian impairment after equalization, and random errors
  - Your mileage may vary
- No AUI errors
- Even with an ideal receiver, the shallow curve for high K means it takes more SNR improvement to get to link with a commercially acceptable frame loss ratio

Bad optical signal, spec-worst AUIs and error floor

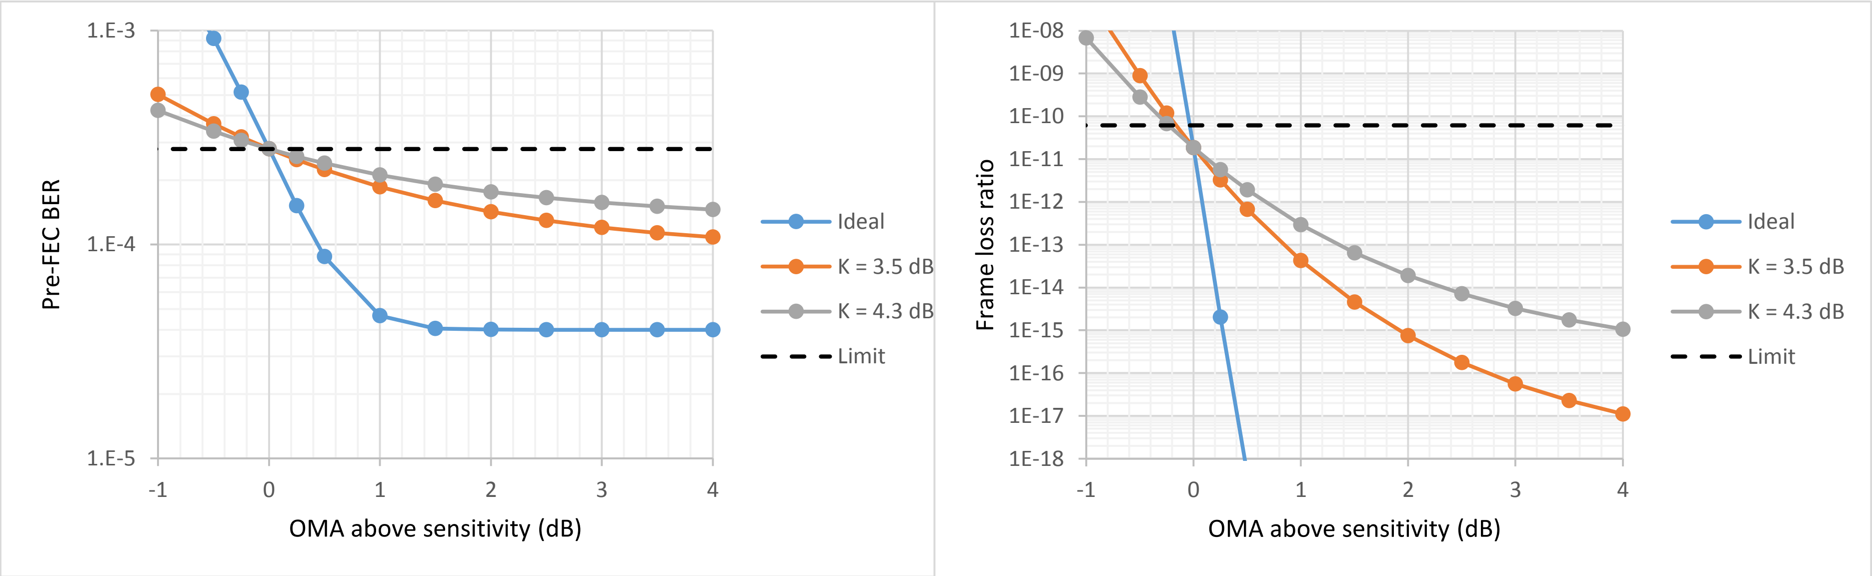

- Textbook theory, assuming Gaussian impairment after equalization, and random errors
  - Your mileage may vary
- Maximum AUI errors
- Even with an ideal receiver, the shallow curve for high K means it takes more SNR improvement to get to link with a commercially acceptable frame loss ratio

Examples of setting up marginal or poor transmitters

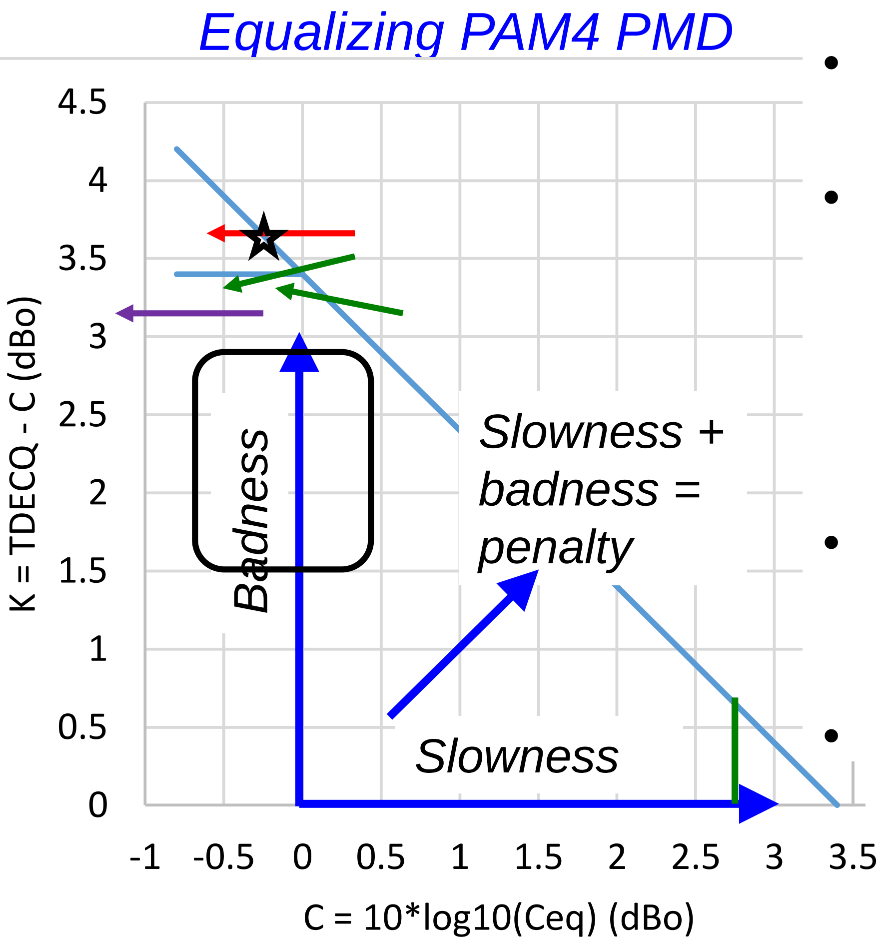

- From the start point of a green or red arrow;
- Changing the transmitter setups can make the two green transmitters pass
  - if they don't fail any overshoot spec
  - If transmitter power or swing is limited, the lower green arrow hasn't improved the link
    - OMA down as emphasis up
- Changing the transmitter setup didn't make this red transmitter pass
  - It shouldn't; it is too noisy or distorted
- The left end of the purple arrow falls foul of the minimum cursor tap weight (so, non-compliant as it should be)

Black box and star show the area of most cases in rodes_3cu_adhoc_030520_v2 (see slide 18) and the very worst case. Worse would be allowed without K or maybe overshoot limit – but should not be allowed

Examples of setting up marginal or poor transmitters

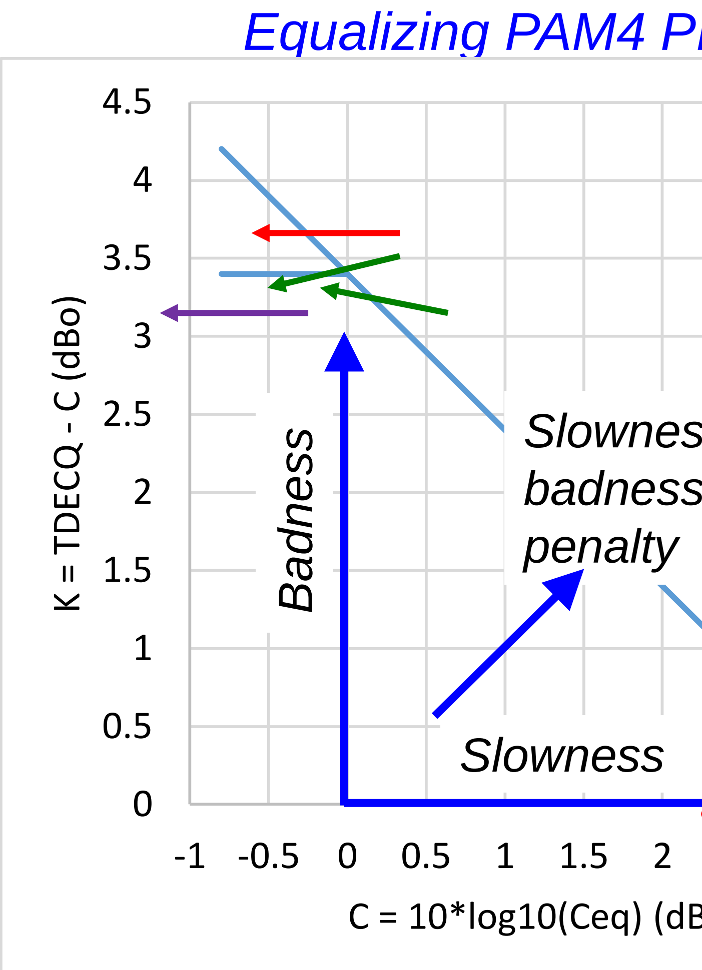

- From the start point of a green or red arrow;
- Changing the transmitter setups can make the two green transmitters pass
  - if they don't fail any overshoot spec
  - If transmitter power or swing is limited, the lower green arrow hasn't improved the link
    - OMA down as emphasis up
- Changing the transmitter setup didn't

TDECQ - $10\log_{10}(C_{EQ})$ Explanation

- TDECQ - $10\log_{10}(C_{EQ}) \leq TDECQ_{MAX}$ *Table 140-6, 151-7*

1. $TDECQ \approx TDECQ_{MAX}$
   - $10\log_{10}(C_{EQ}) \geq 0 \leftrightarrow C_{EQ} \geq 1$
   - $H_{EQ}(f)$ can not adopt to Low Pass
   - TX pre-emphasis qualitatively constrained
2. $TDECQ < TDECQ_{MAX}$
   - $10\log_{10}(C_{EQ})$ unconstrained $\leftrightarrow$ C_{EQ} unconstrained
   - $H_{EQ}(f)$ unconstrained
   - TX pre-emphasis unconstrained

- TDECQ - $10\log_{10}(C_{EQ})$ is just a lousy pre-emphasis spec

Yes, it's that simple!

20 January 2020 7 II-VI

No, you can do as you please (green arrows) as long as there isn't too much noise and distortion (red arrow) or you go too far (purple arrow)

No. K is not a pre-emphasis spec, it's a badness spec. It's that simple.

From
http://ieee802.org/3/cu/public/Jan20/cole_3cu_01b_0120.pdf#page=7

Example real transmitters

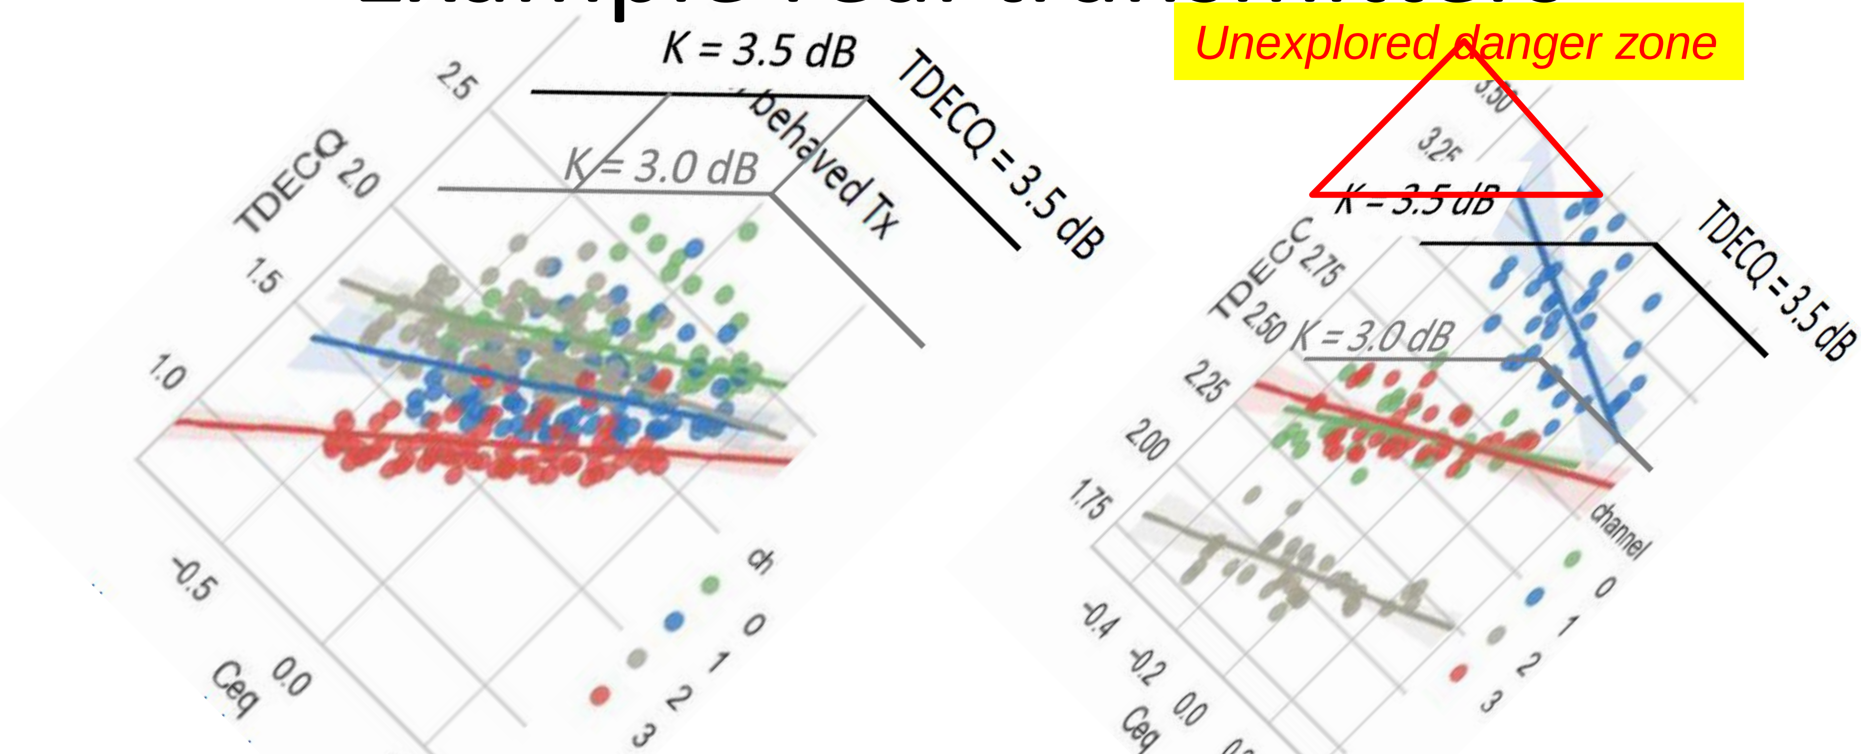

- The four signals on the left and three on the right all pass; they are not significantly affected by the K limit
- One of the four signals on the right passes unless the emphasis is turned up too much
- Spec is working
- Based on http://ieee802.org/3/cu/public/cu_adhoc/cu_archive/rodes_3cu_adhoc_030520_v2.pdf#page=22

Equalizable overshoot again

TDECQ - $10\log_{10}(C_{EQ})$ Explanation Confirmation

- Piers Dawe analysis
  http://www.ieee802.org/3/cd/public/July18/dawe_3cd_01b_0718.pdf#page=11

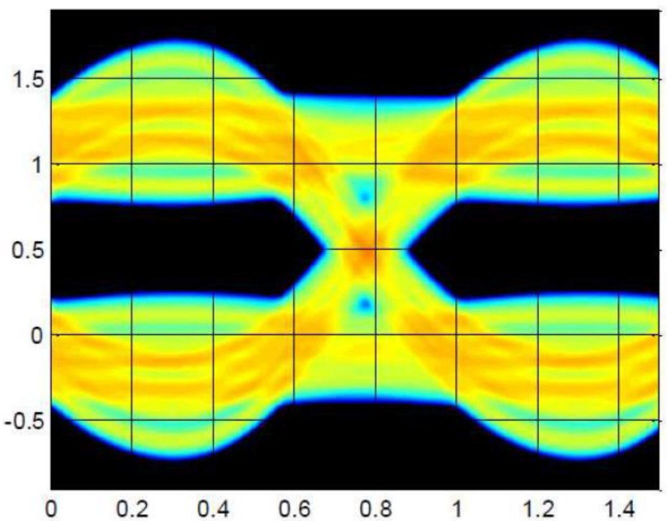

"very bad" TX eye that can be stopped by TDECQ - $10\log_{10}(C_{EQ})$ i.e. "bad" or "very bad" = pre-emphasis

20 January 2020 9 II-VI

From
http://ieee802.org/3/cu/public/Jan20/cole_3cu_01b_0120.pdf#page=9

- This signal would be allowed by the tap weight rules, and excluded by the K = TDECQ – $10\log_{10}(Ceq)$ rule
- "bad = pre-emphasis" is not correct
- In dawe_3cd_01b_0718 and in TDECQ theory, overshoot was/is not bad in itself
- But the thick bands are really bad for PAM4!
  - Remember this is a PAM2 eye
  - K = TDECQ – $10\log_{10}(Ceq)$ screens out this bad signal
  - because it's noisy or distorted
- But for real receivers with real impairments, too much overshoot must be bad

- The minimum cursor tap weight in the TDECQ method limits *equalizable* overshoot at the receiver
- Need this to protect the equalizer, but probably need something more to protect the receiver front end

Including more realistic receiver impairments

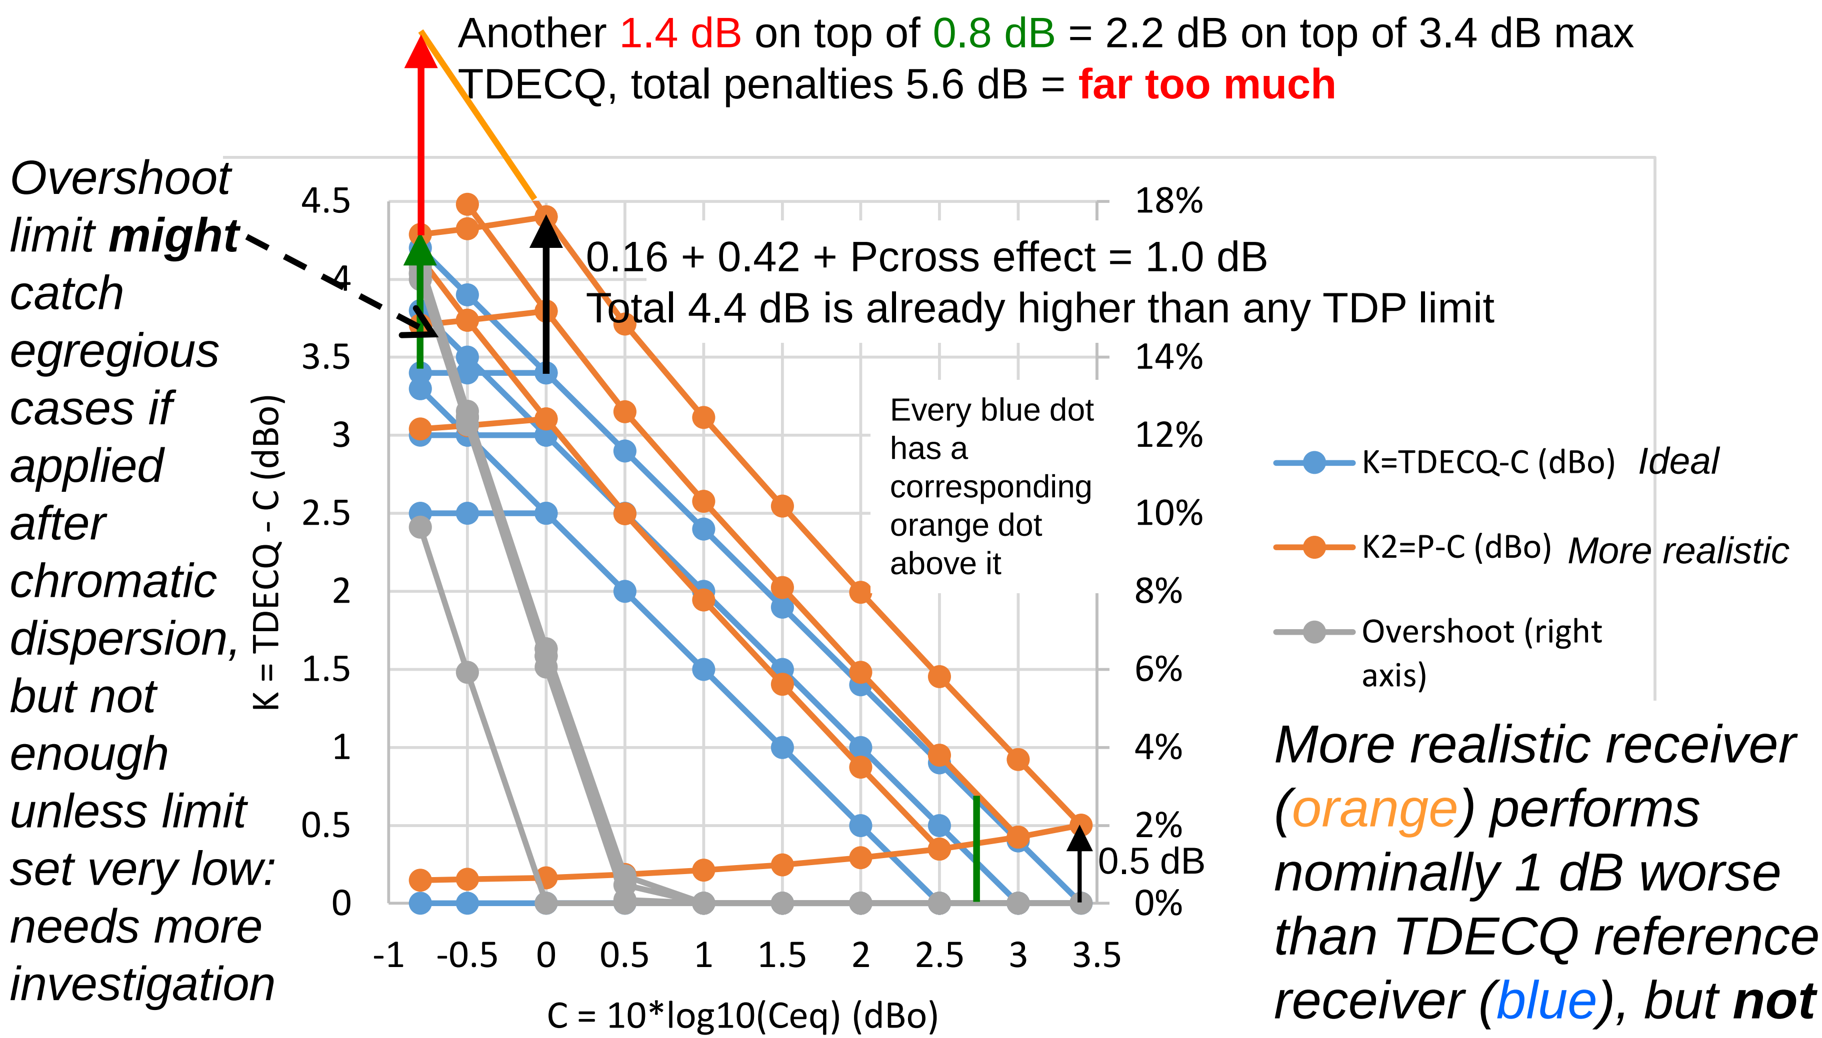

*More realistic receiver (orange) performs nominally 1 dB worse than TDECQ reference receiver (blue), but **not uniformly***

Different impairments affect the right and top areas

Bad optical signal, spec-worst AUIs and realistic receiver

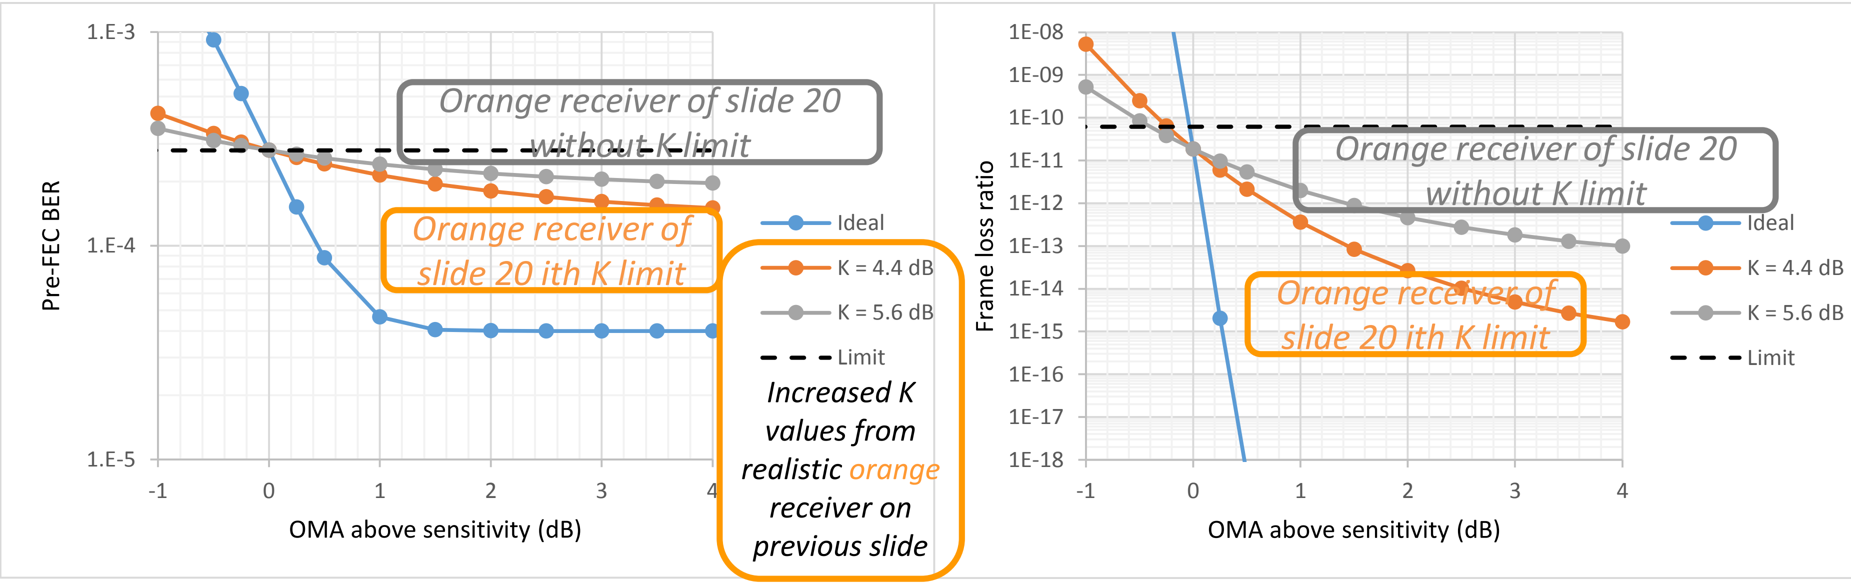

- Textbook theory, assuming Gaussian impairment after equalization, and random errors
- Maximum AUI errors
- Orange receiver from previous slide
- With reasonable receiver impairments, the shallow curve for high K means it takes a lot more SNR improvement to get to link with a commercially acceptable frame loss ratio

The real power budget

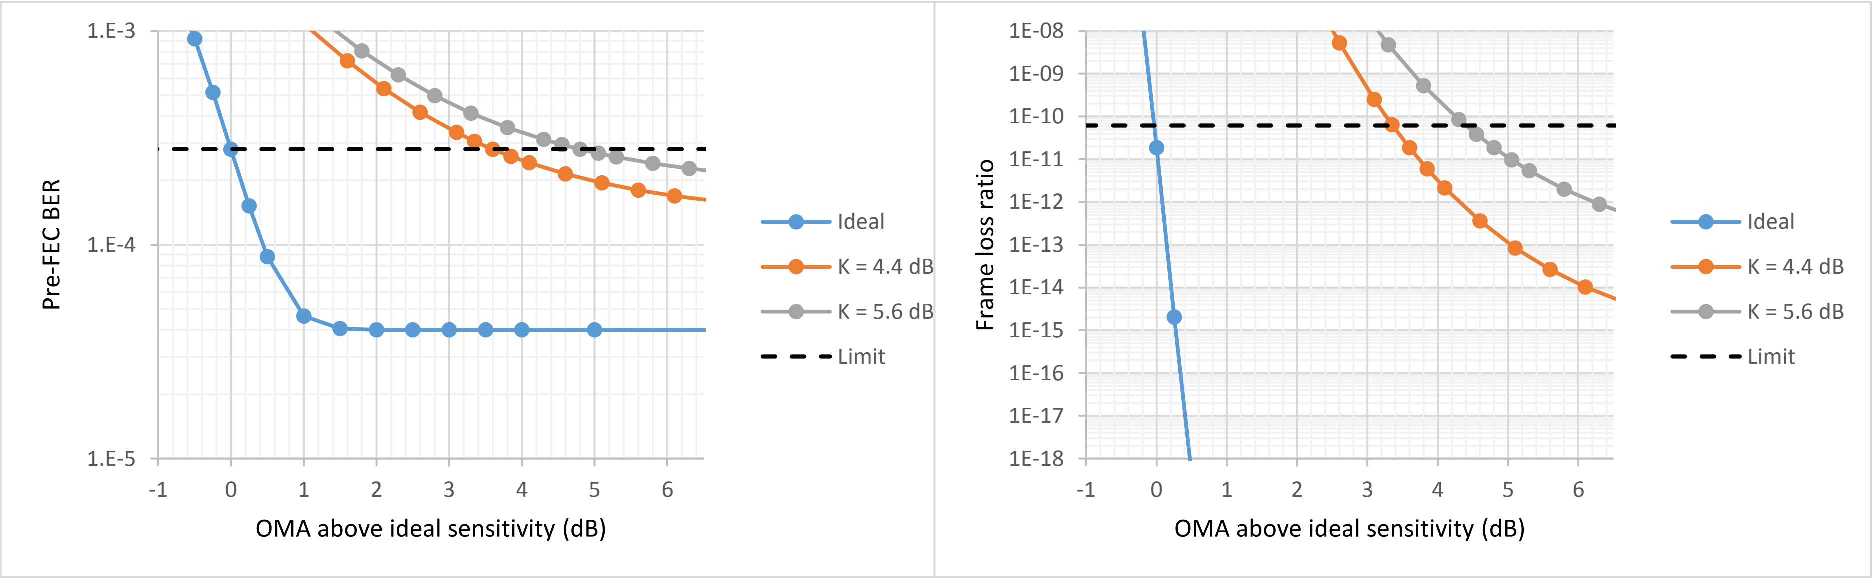

- With reasonable receiver impairments, the shallow curve for high K on top of a high nominal penalty (TDECQ) creates a weak link
- We don't need it
- No other SMF spec would allow it
- Transmitters can avoid the high K corner

Summary of relation between TDECQ and overshoot

- Two sorts of overshoot, equalizable (emphasis) and not
- **Equalizable** overshoot is constrained by minimum of largest tap coefficient (0.8)
  - **Towards the left** on the TDECQ map
- **Unequalizable** overshoot is in the same category as all unequalizable ISI and signal noise, measured by K
  - **Towards the top** of the TDECQ map
- K is like TDP, TDEC and EVM: a necessary screen for bad signals. All modern optical specs have one
  - K was never intended to screen for too much emphasis; largest tap coefficient (from May 2018) already does that

Non-equalizable overshoot, questions

- See http://ieee802.org/3/cu/public/Jan20/cole_3cu_01b_0120.pdf slides 12 to 15
- Should know what the threat mechanism is before we really know how to screen it out
- E.g. the overshoot could cause bad clipping in the receiver
  - Where in the receiver?
  - Is it overshoot relative to signal or peak-to-peak swing that matters?
  - If the former, the overshoot should be measured **after** the chromatic dispersion (i.e. at TP3), as TDECQ is
  - If the latter, the allowed swing at TP3 could be larger than at TP2 by the *minimum* loss of the fibre
- Next slides address measuring total overshoot
  - We have a spec in place for equalizable over-emphasis, probably too loose
- How do we determine how much an overshoot limit should be?

Thoughts on measuring overshoot

- Limiting overshoot at TP2 may not be enough if chromatic dispersion can make the overshoot higher at TP3, as seems very likely
- A measurement on a square wave measures the worst of pre-emphasis and post-emphasis, but a real signal's overshoot can be determined by the sum of these
- Square wave is a bad choice of pattern anyway because PMAs may fail to lock on it and forward the signal correctly to the PMD transmitter output (as 120.5.11.2.4 says)
- Traditional peak measurements are distorted by scope noise, particularly for optical scopes at such high bandwidths

How to measure overshoot 1/2

- Apply the spec to the same cases as TECQ and TDECQ: TP2, TP3 with most positive chromatic dispersion, and TP3 with most positive chromatic dispersion
- Use the same pattern and observation bandwidth as for T(D)ECQ so that determining the overshoot is another free by-product of measuring for T(D)ECQ, with a much simpler, non-iterative, calculation
- Find the scope noise
- Create a vertical histogram from the measured waveform (not the equalized one)
  - Width of histogram is 1 UI (or e.g. 2, as SSPRQ's length is an odd number of UI
- Convolve the histogram with the noise that could be added to it at maximum T(D)ECQ, e.g. OMAouter*$10^{-3.4/10}$/(6*3.414), RSS-reduced by the scope noise

How to measure overshoot 2/2

- Find the points on the two outer tails where the CDFs come to a number such as 5e-5
- Either find the distance from the "three" level to the upper point, and from the lower point to the "zero" (these are the overshoot and undershoot before normalisation)
- or find the distance from the average level to the upper point, and from the lower point to the average (these are the peak excursions)
  - Which one is more relevant to the receiver?
- Normalise by either OMA or standard deviation of the waveform, or don't normalise but scale for minimum TP2-TP3 loss
  - The first is more familiar, the second avoids some error from the pattern dependency of the OMA definition
- Limit upper and lower separately because excursions on just one side could overload a receiver
- To do: choose the limits – based on what information?

Conclusions

- Receivers that cope with unnecessary corner cases cost more power, and typically, something else is compromised
- Receiver protection should not be removed retrospectively
  - But adding an overshoot limit is attractive. Can avoid test cost by re-using TDECQ measurement, but much simpler, quicker calculation
- Limiting C, C+K, and K are all worth doing
  - K just as we included a TDP limit in previous standards
    - I had offered a more gentle limit that slopes up to the left in 2018, but it got no traction
  - Protect against transmitter evading the intention of the spec
- Limiting K is as practical and convenient as TDECQ, a by-product of TDECQ, and what we have been used to since July 2018 (P802.3cd/D3.4)
  - No Tx measurement cost, avoids extra Rx measurement cost
- Stay with K limit, add spec for overshoot at TP3 as well as TP2