Stressed receiver sensitivity setup and calibration
(comments R1-304 R2-196, R2-203)

Piers Dawe and Jakob Toft

NVIDIA

Supporters

- John Calvin Keysight
- Steen Christensen NVIDIA
- Johan Mohr NVIDIA

Introduction

- The signal for stressed receiver sensitivity is defined in 180.9.15
- Following 121.8.10: a clean signal is degraded with sinusoidal jitter, sinusoidal amplitude interferer (SI), Gaussian noise (GN) and a low-pass filter
- There are some constraints (changed in D3.1) and the implementer has quite a lot of freedom
- Different implementations can give different results, and unrealistic levels and types of stress
- SI and its calibration is a particular problem

Figure 121–8—Stressed receiver conformance test block diagram

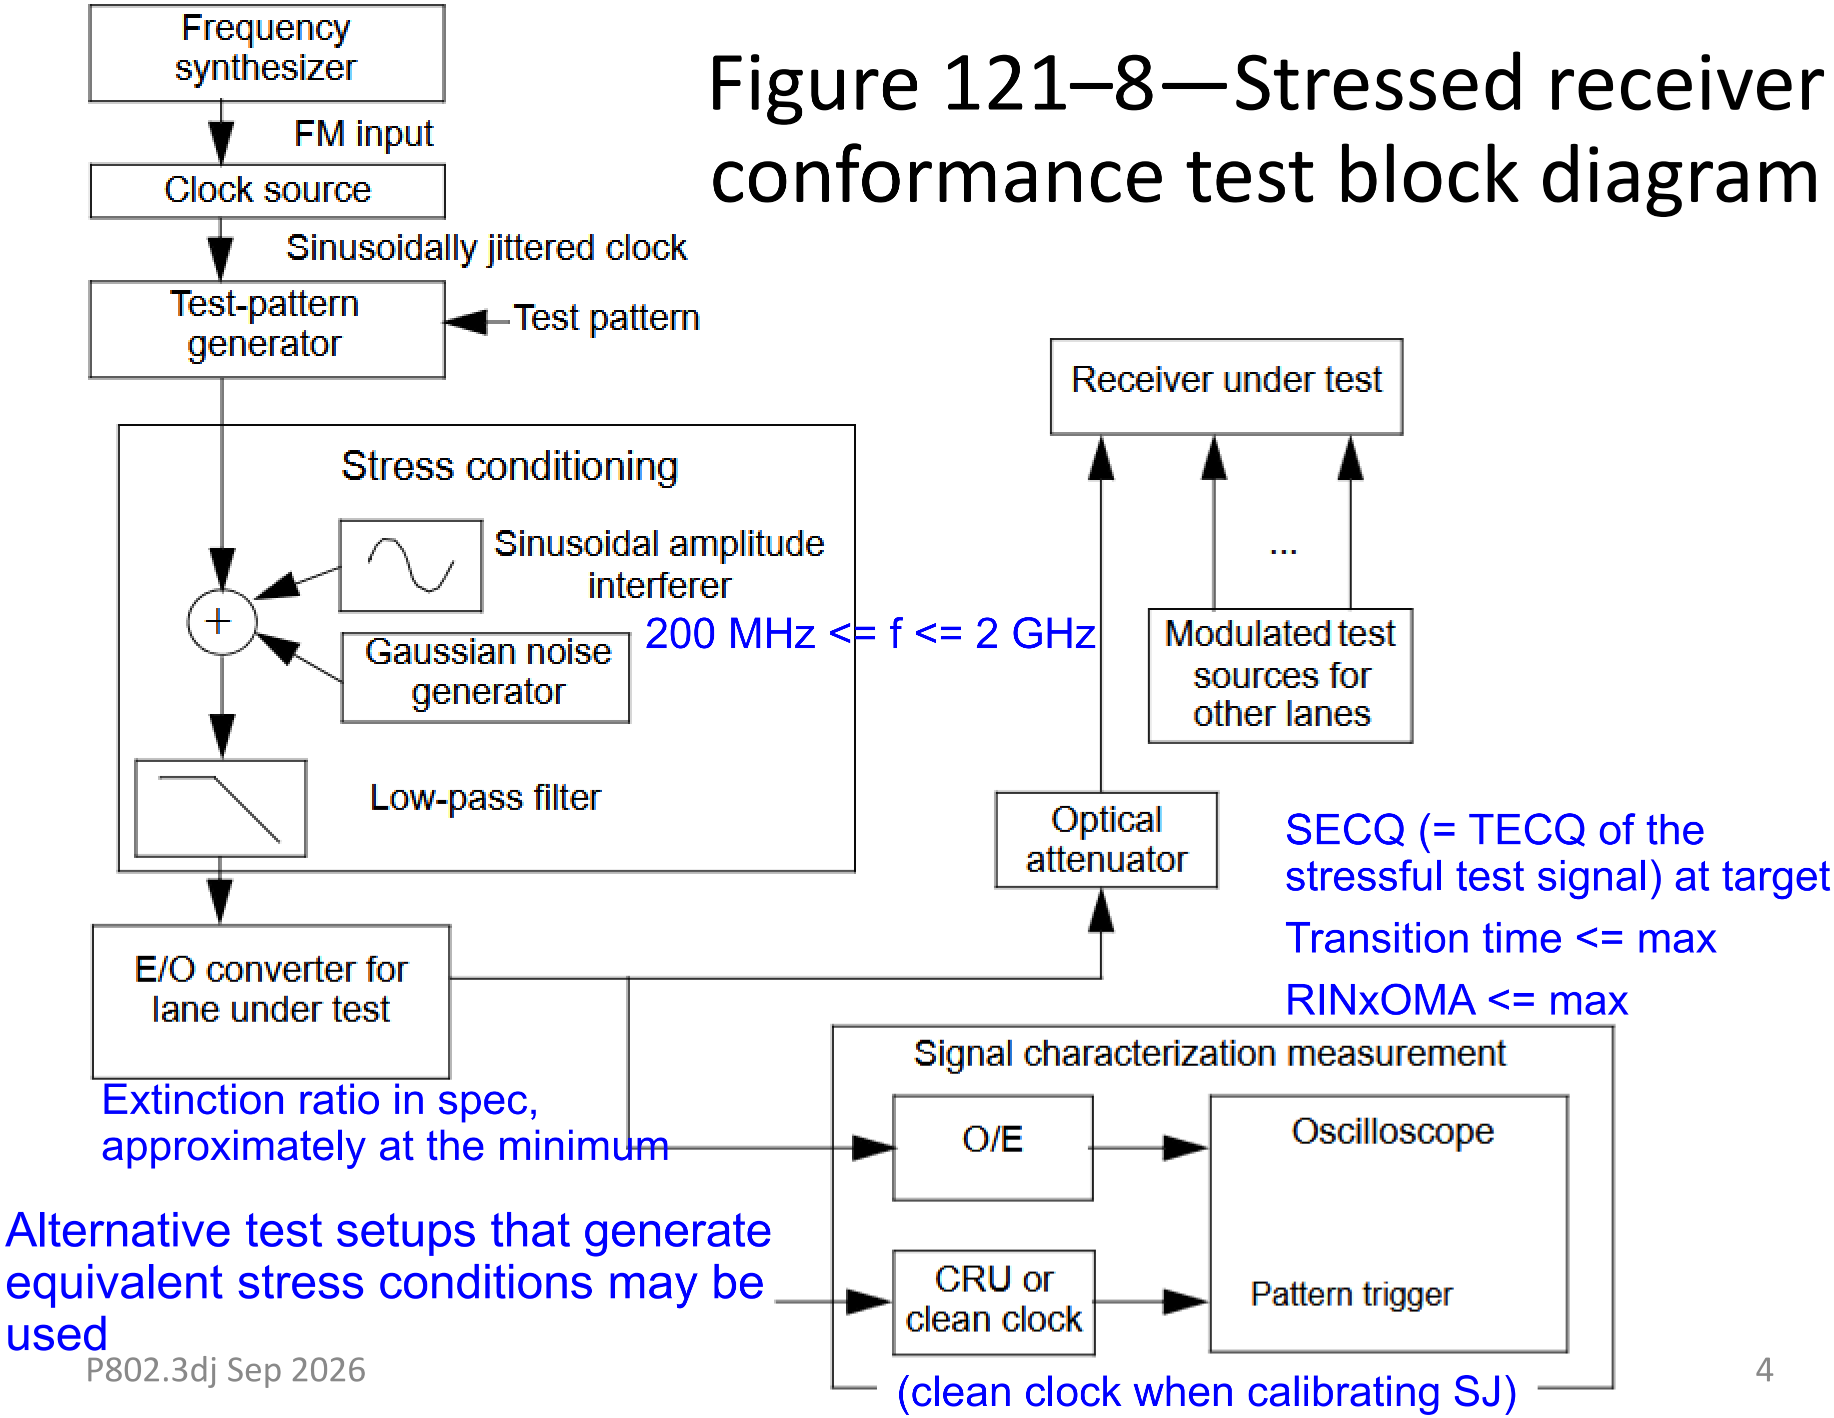

From 121.8.10.1

- The sinusoidal amplitude interferer may be set at any frequency between [200] MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate. The Gaussian noise generator, the amplitude of the sinusoidal interferer, and the low-pass filter are adjusted so that the specified SECQ is met.
- For improved visibility for calibration, all elements in the signal path ... should have wide and smooth frequency response, and linear phase response, throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be negligible.

The test implementer chooses the mix of filtering, Gaussian noise and SI amplitude. Also the SI frequency

Some choices can give very unrealistic signals

Often the penalty from SI is very large

Comments R1-304 and R2-196

- R1-304 SC 180.9.15 P 493 L 45 TR
- The frequency of the sinusoidal interferer was changed from minimum 100 MHz to minimum 200 MHz, but it's still a single frequency. The sampling scope aliases it to some other frequency and the software FFE will amplify it according to the apparent frequency, not the real one, **causing error**.
- 1. Change from a sine tone to a PRBS generator (not synchronous to the victim) with a high enough signalling rate that wherever the aliasing takes the apparent frequencies to (there are many real frequencies in a PRBS), the overall enhancement will be reasonably consistent. A PAM2 PRBS is more "bounded-like" than a PAM4 one as well as having a uniform spectrum.
- 2. Limit the amount of bounded interference.
- 3. In 121.8.10.1, we say "care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate". After taking advice from scope experts: add the sampling rate of the scope to that list.

- R2-196 SC 180.9.15 P 494 L 4 TR
- 2GHz tone for making the transmitter used in compliance testing is too restictive for many receivers. Noise is expected to be spread across the full spectrum. It should not be expected to have such transmitters in systems.

Comment R2-203

- R2-203 SC 180.9.15 P 510 L 28 TR
- The frequency of the sinusoidal interferer was changed from minimum 100 MHz to minimum 200 MHz, but it's still a single frequency (comment R1-304). The minimum* remains at 2 GHz, same as it was at 10 Gb/s. Back then, for PAM2, without equalisers or FEC, that was a range from 4% to 39% of Nyquist; now it's 0.4% to 1.9%. Most unequalisable distortion, and reflections and crosstalk have spectra like the victim or weighted to higher frequencies, not lower. Also a stressed signal with significant sinusoidal interferer can have bimodal levels (which didn't really matter for PAM2 without equaliser). So, both the spectrum and the distribution are unrealistic and not suitable to subject real receivers to. The sampling scope aliases the real SI frequency to some other frequency which can be much higher than Nyquist and the software FFE will amplify it according to the apparent frequency, not the real one, causing a calibration error which can be surprisingly large. A PAM2 interferer has a better spectrum than the sinusoid but the same bimodal issue.
- 1. Change from a sine tone to a QPRBS generator (not synchronous to the victim) with a signalling rate similar to the victim and passed through the same or similar filtering to the victim.
- 2. Apply a maximum limit to the bounded interference (now QPRBS), or apply a minimum limit from filtering, e.g. minimum 1.9 dB from filtering alone.
- 3. In 121.8.10.1, we say "care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, and the pattern repetition rate". Modify this as necessary for the different interferer and after taking advice from scope experts, describe how the frequency of the interferer and the sampling rate of the scope should be related.

* *Should have been maximum*

Error in the SRS measurement – spectrum

- In general, we expect that a signal-borne impairment such as distortion or the effects of reflections has a spectrum related to the "victim".
  - Distortion — Similar spectrum to victim
  - Crosstalk — Victim's spectrum, strongly weighted to higher frequencies
  - Reflections (uncorrected) — Victim's spectrum, usually weighted to higher frequencies, or could be unweighted
  - Power supply noise — Few MHz or lower
- We have one interferer and the Gaussian noise to represent the combination of all of these
  - By the way, that combination will add to the Gaussian skirts (distortion and reflections are excluded from the RIN measurement)
- PSU noise should not dominate
  - If we are very concerned about it, that would be a separate tolerance spec with a MUCH lower interferer frequency
- The single SI frequency, which the receiver must tolerate, is unrealistic
- The strong single-tone nature of the SI could confuse some repetitive process in a DSP

Error in the calibration

- The aliased frequency(s) could be anything
  - It could be a single tone, anywhere from DC to signalling rate × oversampling rate / 2 (that's 1700 GHz!)
  - Or it could be aliased to many frequencies
  - The aliasing takes place after the scope front-end filter, so the software FFE can amplify it
  - Not what happens in a real receiver
- The scope that aliases to many frequencies calculates noise enhancement (Ceq) on the basis of its front-end bandwidth (fb/2) when it calculates SECQ; for SI, it estimates an FFE gain that isn't really there
- By comparison, the Gaussian noise spectrum can be flat up to a cut-off at the implementer's choice up to 64 GHz
  - which is above Nyquist
- But the SI is several times larger than the Gaussian noise

Unrealistic bimodal distribution

- A sinusoidal interferer is close to "dual-Dirac"
- PAM2 PRBS has the same problem
- Because the interferer is quite strong, the way people build stressed signals today, the resulting PAM4 histograms have two peaks for each level
- This is unrealistic

Before adding Gaussian noise

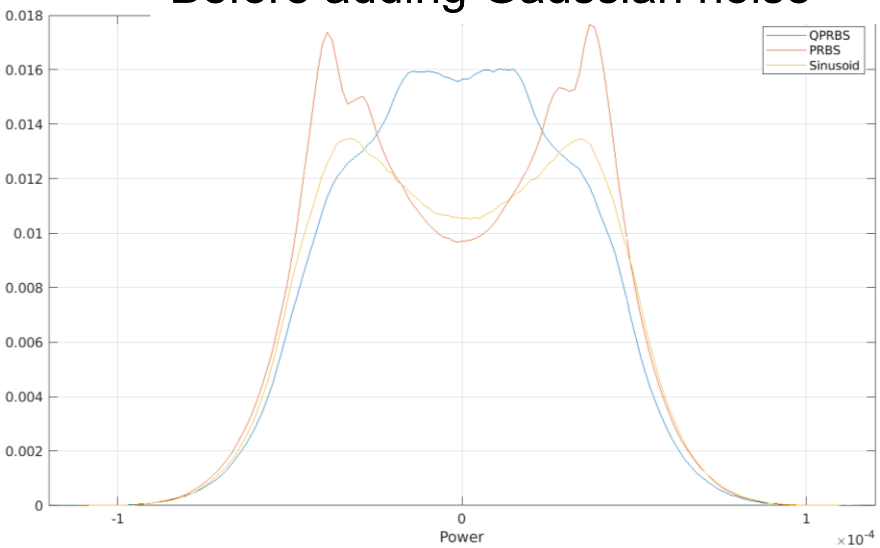

After adding Gaussian noise

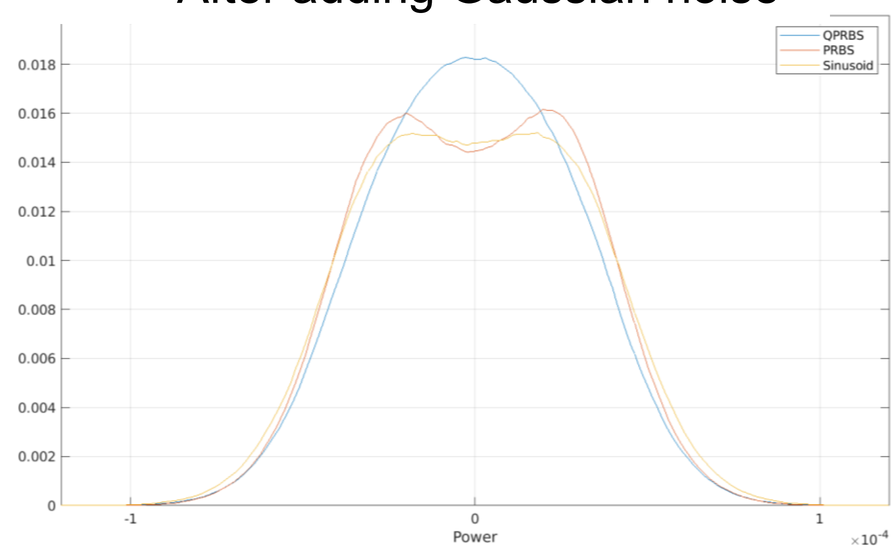

Simulated stressful signal for SRS testing: vertical pdfs of 4 levels combined

History lesson

- The SI maximum frequency is very old: it was set at signaling rate / 5 (then 10.3125 GBd) before equalizers were expected:

52.9.9.1 Stressed receiver conformance test block diagram

The frequency of the sinusoidal interference may be set at any frequency between 100 MHz and 2 GHz, although be careful to avoid a harmonic relationship between the sinusoidal interference, the sinusoidal jitter, the data rate and the pattern repetition rate.

The Bessel-Thomson filter and the E/O converter should have the appropriate frequency response to result in the appropriate level of initial ISI eye closure* before the sinusoidal terms are added.

58.7.11.2 Stressed receiver conformance test signal characteristics and calibration

To emulate the effects of DCD or data-dependent jitter, at least 0.05 but no more than 0.15 UI peak-peak of pulse shrinkage jitter should have been achieved. This imposes a limit of less than 1.2 dB of vertical closure from sinusoidal interference, applied after vertical closure created by filtering.

The frequency of the sinusoidal interference may be set at any frequency between B / 100 and B / 5 where B is the signaling speed†, although care should be taken to avoid a harmonic relationship between the sinusoidal interference, the sinusoidal jitter, the signaling speed and the pattern repetition rate.

* e.g. 2.2 dB for 10GBASE-L where max TDP (predecessor to ECQ) is 3.2 dB

† 125 MBd or 1.25 GBd in this case

The SI frequency was comfortably above the CDR bandwidth and low enough so as not to cause pulse width shrinkage

The SI amplitude was capped

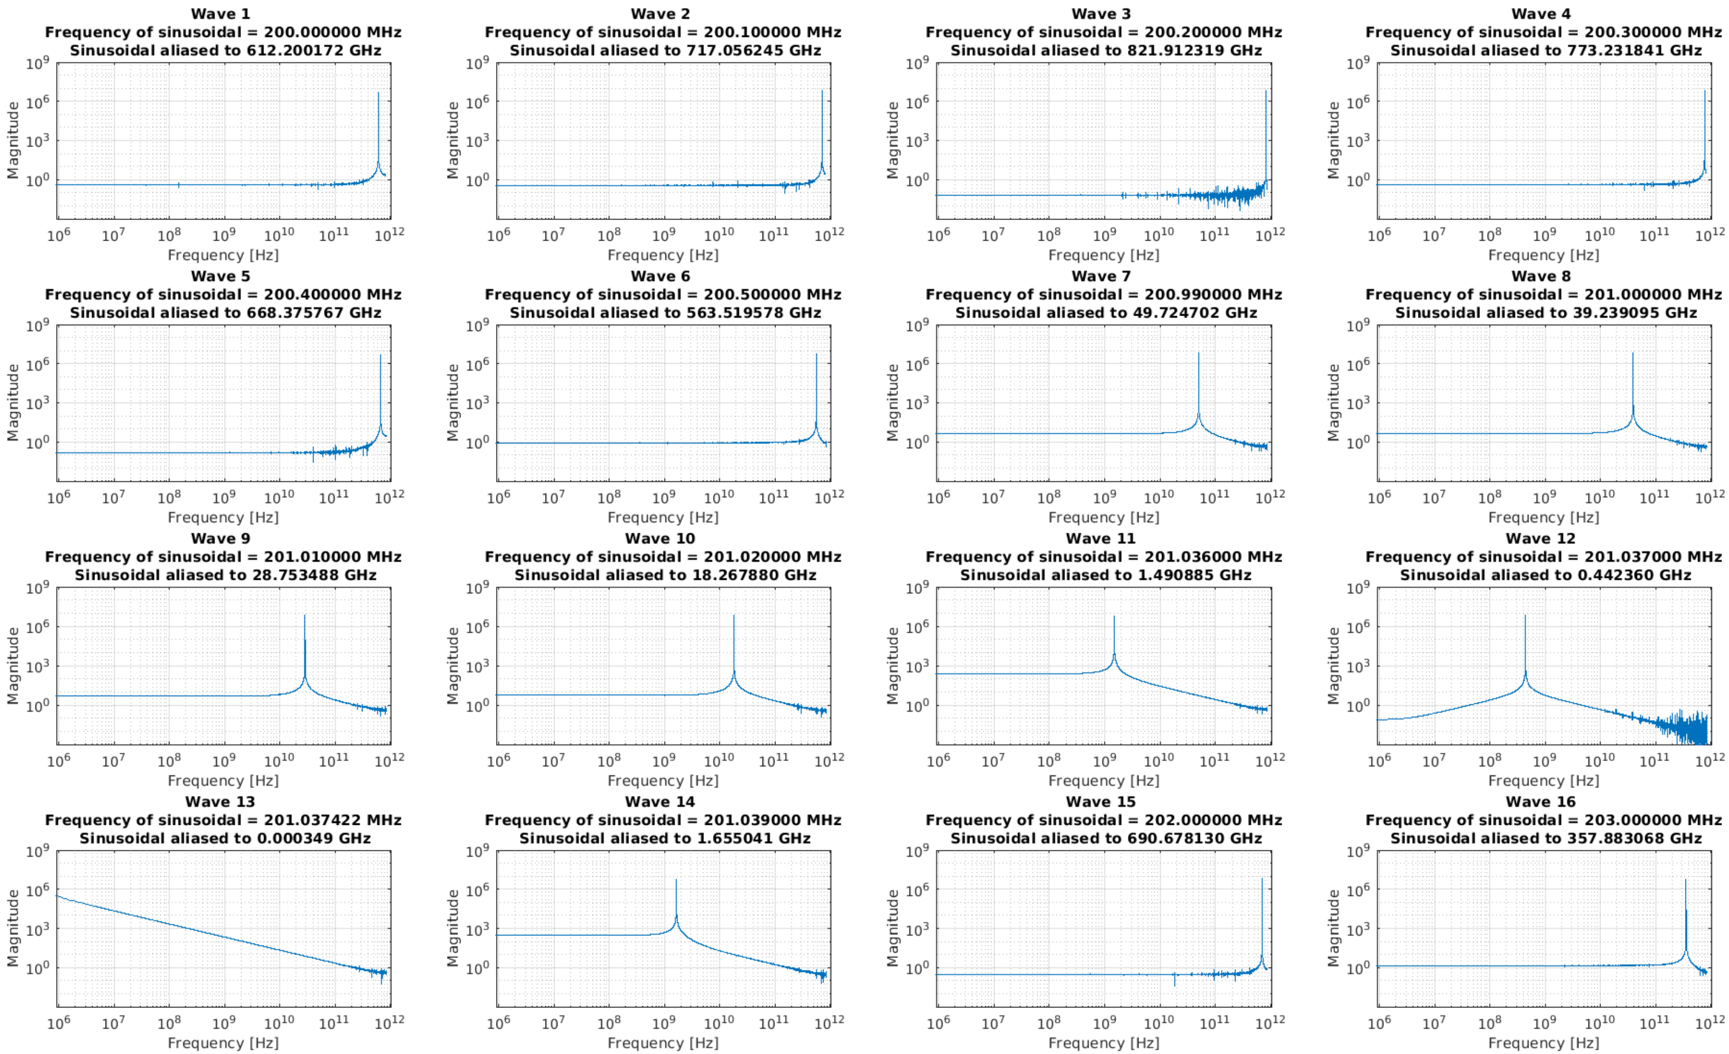

- Examples of aliased SI spectra
  - if the aliasing does not distribute the spectrum
- A difference of 100 kHz in sinusoidal frequency can move the aliased frequency by 100 GHz

SECQ variability

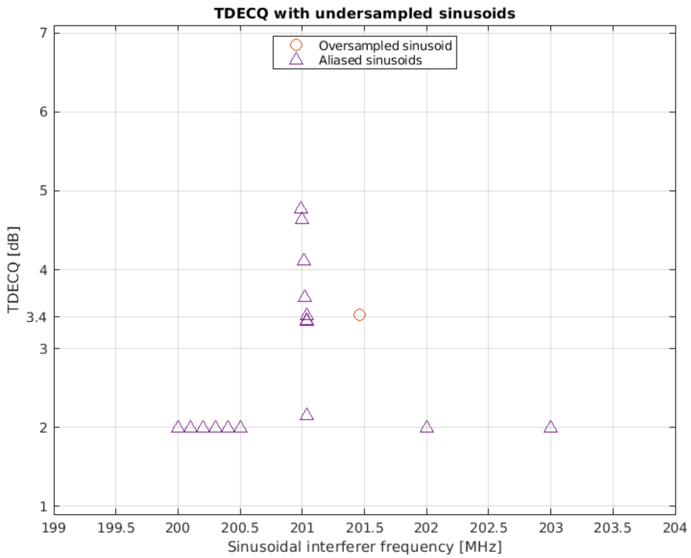

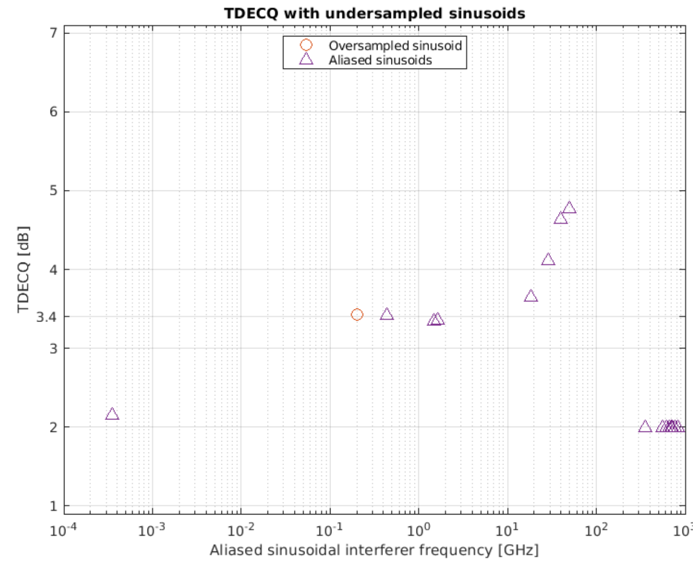

- Simulations with constant SI amplitude, different frequencies as on previous slide. Assuming the aliasing does not distribute the spectrum. Left: SECQ vs. actual frequency; right: vs. aliased frequency
  - Simulations with the same SI amplitude at different frequencies give very different (apparent) SECQ: +/-3.4 dB
- If the aliasing does distribute the spectrum
  - the error is smaller than the extreme, and consistent
- In calibration, the SI amplitude would be adjusted to bring the apparent SECQ to the target, giving very different real stress

Way forward

- With reference to Figure 121–8: change from sinusoidal interferer to PRBS13Q, summed with the others before the low-pass filter as today
- With a QPRBS interferer, there are many frequencies. They are aliased too, but the noise enhancement of the aliased spectrum and the real spectrum is similar so that's OK
- To avoid the 2-moded histogram problem, use a PAM4 QPRBS rather than a PAM2 PRBS
  - PRBS13Q is a well-known pattern with more than enough lines for this purpose but short enough to be generated by an arbitrary waveform generator
- The amount of bounded interferer (SI or PRBS) should be limited
  - Target approximately 2.5 dB from the filtering
  - Before SJ, SI or GN are added. Measured with DFE = 0 for consistency
- More use made of the Gaussian noise generator
- The GN should represent the tails of the distortion/crosstalk/uncorrected reflections in a bad signal as well as worst-case RIN, therefore it should be a little more than the product RIN limit: propose 2 dB above, = -137 dB/Hz
- Recommend that with bounded interferer off and other impairments on, the SECQ should be within 0.5 dB of the target SECQ with all impairments on
  - This leaves the implementer with enough freedom to set the mix

Convenient stop-gap alternative

- If a PRBS13Q interferer is not available, a pair of sinusoidal interferers could be a substitute, for a small measurement error
- The amplitude of the higher frequency signal, after filtering as in Figure 121–8, is half that of the lower frequency
- Suggested frequencies are chosen such that the amplification of these signals by a 2-tap FFE is similar to the noise amplification of the reference receiver front-end frequency response
- Not exactly $f_b/8$ (13.28125 GHz) and $3f_b/8$ (53.125 GHz), and not harmonically related to each other, the sinusoidal jitter, the signaling rate, or the pattern repetition rate

Harmonic relationships

- We use QPRBS aggressors in many places in the standard
  - Using the same rate as the victim is convenient, (whether they are locked together or not)
  - Using a different-length sequence (to the victim) is seen as acceptable for making a co-propagating aggressor uncorrelated
  - Unlocked is OK. If they are locked, the rates should be offset
    - Enough that the interferer UI-level phase walks across the victim in the duration of the signal calibration or receiver stressed sensitivity measurement, not so much that the spectrum is grossly different
  - A longer pattern such as PRBS31Q contains many tones, as desired. Probably PRBS15Q would work, maybe PRBS13Q if it has a practical advantage
- It is already stated that "Care should be taken to avoid harmonic relationships between the sinusoidal interferer, the sinusoidal jitter, the signaling rate, the pattern repetition rate". It appears that it is not necessary to mention the oscilloscope sampling rate. "Sinusoidal interferer" changes

Specific changes

Conclusion