Considering Alien Noise

Adee Ran, Intel

Preface

- With high loss channels (accepted objectives) the noise budget should be carefully inspected.
- We consider ISI and Xtalk, and we specify jitter, but these are not the only sources of noise.
  - Even thermal noise is not negligible any more…
  - What else should we look at?

The problem

- Define any signal that can be felt by the RX when all 8 lanes (thru, NEXT, FEXT) are quiet as “alien noise”
  - That could be Xtalk from other links, or any other EMI present in the system
- How much (or how little) alien noise should we expect?
- Different people have different assumptions… resulting in different results
  - Action item from January meeting: “correlate results with NRZ contributors to close the gaps”.

Updated objectives

- Define a 4 lane PHY for operation over a printed circuit board backplane with a total channel insertion loss of <= 35 dB at 12.9 GHz
- Define a 4 lane PHY for operation over a printed circuit board backplane with a total channel insertion loss of <= 33 dB at 7.0 GHz

How much signal is left?

- Sample channel from IBM has ~35 dB loss at 12.9 GHz; for this channel, COM analysis shows available signal of 17.4 mV (after equalization)
  - Assume allowance of 8 dB for jitter and RX implementation; so total noise budget is ~ 7 mV (peak)
  - More details in ran_01_0712
- We don't have any contributed channel with 33 dB loss at 7 GHz (second objective); but assuming similar equalization, and 1/3 factor for PAM4, the available signal is ~ 6 mV
  - Noise budget is ~ 2.3 mV (peak)…

To illustrate the challenge, consider thermal noise (Johnson model)

- For NRZ:
  - With assumed receiver BW of 19 GHz (4th order filter), thermal noise from channel termination is Gaussian with σ=0.2 mV
  - Assuming FEC, max raw BER is 1e-5, so Q≈4.3; thermal noise "peak" is almost 1 mV.
  - **Thermal noise consumes 13% of the noise budget!**
- For PAM4:
  - BW is 10 GHz, σ=0.15 mV, **peak is 0.7 mV – 28% of the noise budget!**

Can alien noise also reach comparable levels?

- Good design should minimize coupling of external noise sources to the channel
  - Differential pairs, stripline, quality connectors, back-drilled vias…
- An experiment conducted at Intel showed that on a reasonably-designed backplane channel, getting ~ 1 mV of differential signal from EMI is not very likely
  - Requires high-efficiency radiator (e.g. tuned antenna) located within a few cm from the RX package, with >10 dBm fed into it (unlikely to occur from high-speed signaling)
- The RX can still be impacted by common-mode noise, which is harder to control.

So it's not a concern?

- Strong alien noises may be eliminated with reasonable design and common operating conditions.
- However there is currently no written specification that guarantees it:
  - Someone may need to put that antenna near the RX in some application
  - Or operate just below a strong transmitter (e.g. cellular network)
  - Or maybe use non-coupled pair (sometimes it works!)
- Without some requirement, people will expect it to work anyway.

Possible remedies

- Add a specific system/environmental requirement
  - New burden on design
  - Do we have experts that can define tests that should be conducted? Are users familiar with such methods?
  - Doesn't look attractive
- Use the receiver spec as an implicit limit
  - E.g. add some differential mode and common mode noise sources that represent EMI the RX should tolerate
  - Other standards do it (e.g. PCI express)
  - Makes an implicit assumption that the noise in a deployed system won't be higher.
  - Is it good enough?
- Do nothing
  - Maybe it's not a concern?
- Other options?

What's next?

- Start a discussion
- Call for presentations…

References

- Equivalent Noise Bandwidth, Tech note, Tim J Sobering, SDE consulting
- Thermal noise calculator & Formulae, Radio-Electronics.com