10GBASE-T Uncoded Bit Error Challenges & Ramifications

Ramin Farjadrad

Aquantia Corp.

May 12, 2015

Background

- A major challenge of 10GBASE-T in early days was the uncoded bit errors
  - The main cause was a combination of transmitted signal non-linearity and channel limited return loss up to 500MHz
    - Standard spec for worst-case cable RL too relax
  - The above combination resulted in large non-linear echo signal that PHY linear filters could not cancel
  - The non-linear echoes were data dependent and for certain data patterns, they could add up constructively into peaks larger than receive uncoded bits hamming distance, despite their 12dB SNR margin advantage!
  - Although probability of occurrence for such event was very low, but its rate only had to be >1E-12

Channel Example & Reflections

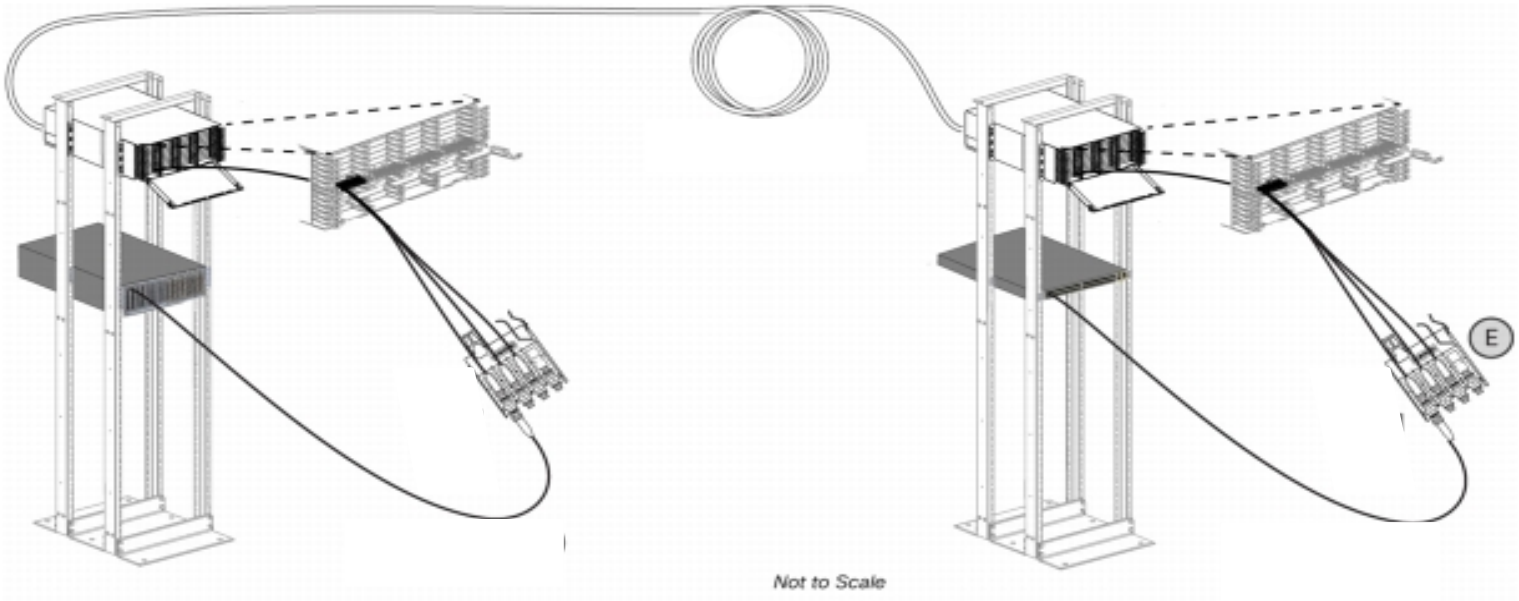

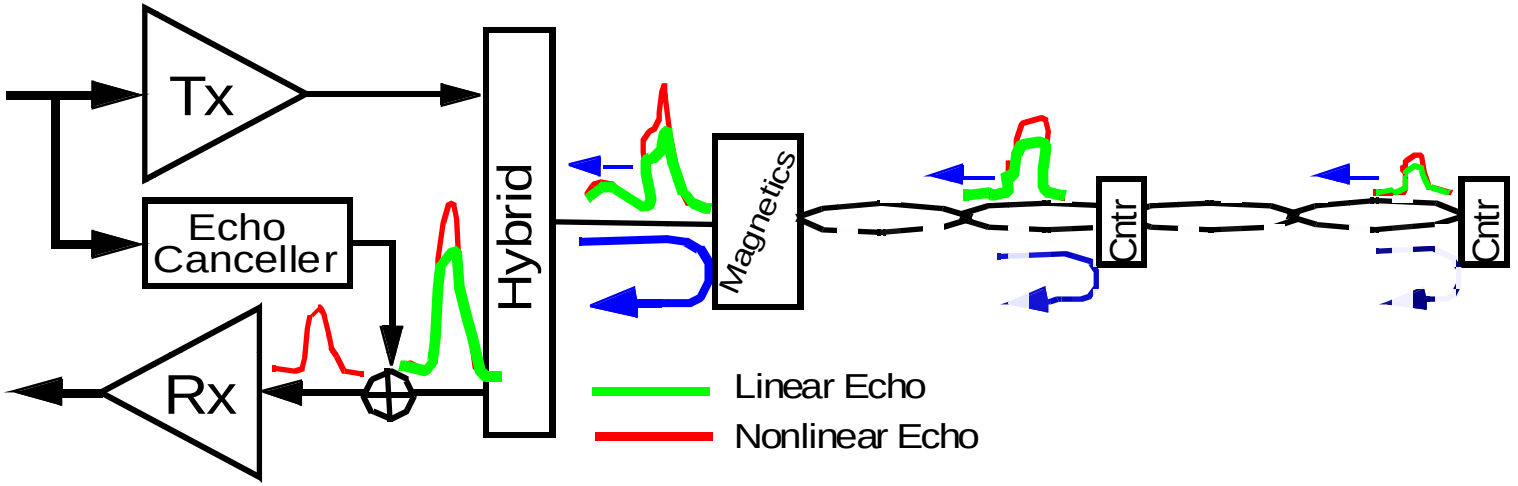

Industry Solutions to The Problem

- Industry had to address both problem sources
  - PHY Vendors
    - Ultra-Linear Tx Driver → Complex & High Power
    - Slew-rate Controlled Driver
      - High Capacitor at Tx Output → Bad RL
      - Oversampled/Closed-loop Opamp Driver→ High Power

    ➔Driver Power in AFE >1W ➔ Better Cooling ➔ High Cost

    ➔Tx Driver THD > 60dB ➔ Long design cycle➔ Late Products
  - Magnetic/ICM Manufacturers
    - Improved RL up to 400MHz→ Challenge→ High Cost
  - System OEMs
    - Stricter rules for MDI traces→ Complex board design→ Higher Cost
    - Sensitive to damaged UTP/Connectors → Field Returns

A Big Cost for a Small Problem!

- Industry paid a high cost only to eliminate very rare single error events occurring less than once every 10billion bits!
  - This happened when link had solid average SNR margin
- Protecting the uncoded bits against but large noise peaks could save industry significant time and money
  - Tx linearity target was the biggest challenge of 10GBASE-T
    - Long design cycles & High power → Delay in viable products
  - Each PHY vendor needs custom magnetics matched to its Tx
    - The cost of 10GBASE-T ICM per port close to cost of a PHY!
  - Transmit design challenges, High power and High Cost
    - Late adoption by the industry

A Better Solution?

- An FEC extended to uncoded bits removes errors caused by rare but large noise spikes.
  - Relaxes the PHY transmit linearity spec significantly
    - Lower power + Easier design
      → Lower cost + Earlier availability
  - Relaxes the magnetic RL requirement
    - No custom ICM per vendor→ Lower Cost + Wider availability
  - Makes life easier for the system OEMs
    - Wider availability of PHY/Magnetic vendors
    - Less complex board designs and components
    - Less field returns due to marginal cable specs

The Enterprise Scenario

- Data Centers are built cleanly from scratch for a next generation and with a new cabling

- Enterprise networks want to use and extend the life of the existing cabling as much as possible
- Many of such cabling setups are old and not carefully setup.

An Optimum Solution for Enterprise?

- Enterprise cabling using Cat5E/Cat6 is much more unpredicted compared to data centers
  - Data center cabling setup is done with 10GBASE-T spec in mind
  - Enterprise cabling was done with 1GBASE-T in mind (0-62MHz)
    - 2.5GBASE-T/5GBASE-T with a 100MHz/200MHz wide spectrum will be sensitive to channel anomalies above 62MHz that most enterprise cablings may have ignored
    - Bad return loss due to cable damage/defects show up at higher frequency, which are don't-care at 1GE, but matter for 2.5G/5G
- Enterprise PHY solutions must be cheaper for adoption
  - PHY & ICM prices at high $$/port (as in 10G) not acceptable
- There is a strong need for a FEC on all symbols, driven by multiple economic and time-to-market factors

Conclusion

- Nonlinear echoes from channel will add up and lead to large peaks that are not cancelled by linear echo filters
- Unless PHY transmitter is ultra linear and/or channel RL is high up to 400MHz, these peaks exceed uncoded bits Euclidean distance
  - Both of the above lead to PHY design complexity, Custom magnetics, and potential field issues
  - Such concerns will be worse for 2.5GE/5GE in enterprise where cabling setups were done for 1GE
- A fully coded scheme eliminates such concerns
  - Leads to lower cost & complexity in PHY, ICMs, System design
  - Minimizes field issues at unknown enterprise environments