A PAM2 based solution for lower cost (KR) backplane channels

IEEE P802.3bj 100 Gb/s Backplane and
Copper Copper Cable Task Force Cable Task Force

Sanjay Kasturia
Inphi/PLX Technology
Hamid Rategh
Inphi
Nima Ferdosi
PLX Technology

Atlanta, GA, USA, Nov 8

Synopsis

- Several presentations indicate that NRZ (PAM2) will achieve required error rate with 25Gbs/lane on a pristine channel composed of low loss PCBs (e.g. Megtron 6) however PAM2 at ~25Gbaud will provide no coverage on legacy (or even soon to be launched) 40G backplane channels
- Other presentations show 100G is feasible with wide (but not 100%) coverage on legacy backplane channels using PAM4
  - Complexity of implementation of PAM4, specially a digital implementation, is of concern
  - A digital implementation with 3tap TX de-emphasis, 32 tap FFE, 2 tap DFE, CTF, ADC & FEC (parthasarathy_01_0511.pdf) was shown to cover most legacy channels
  - dabiri_01_1111.pdf shows PAM4 with precoding can address certain demanding channels
- The purpose of this presentation is to draw attention to that fact that PAM4 is not the only way to support legacy channels and that it is possible to support legacy channels with PAM2
  - This does not claim that PAM4 complexity is unmanageable

100G with PAM2 on legacy channels

- Operating on legacy channels requires us to keep the signal bandwidth comparable to that of 10GBASE-KR as the IL spec allows steep roll-off after 6GHz
- Keeping the signal bandwidth close to ~6GHz limits the baud rate to <~12Gbaud
  - PAM 4 proposals get to ~25Gb/s per lane by going from 1 bit/symbol to 2bits/symbol
  - The other way to get to 100G is to double the effective number of lanes
- Doubling the effective number of lanes is possible without doubling of the number of physical lanes because each legacy 40G link consists of 4 unidirectional lanes in each direction
  - By utilizing each lane as a full duplex channel, we get 8 effective lanes in each direction

Block diagram of full-duplex AFE

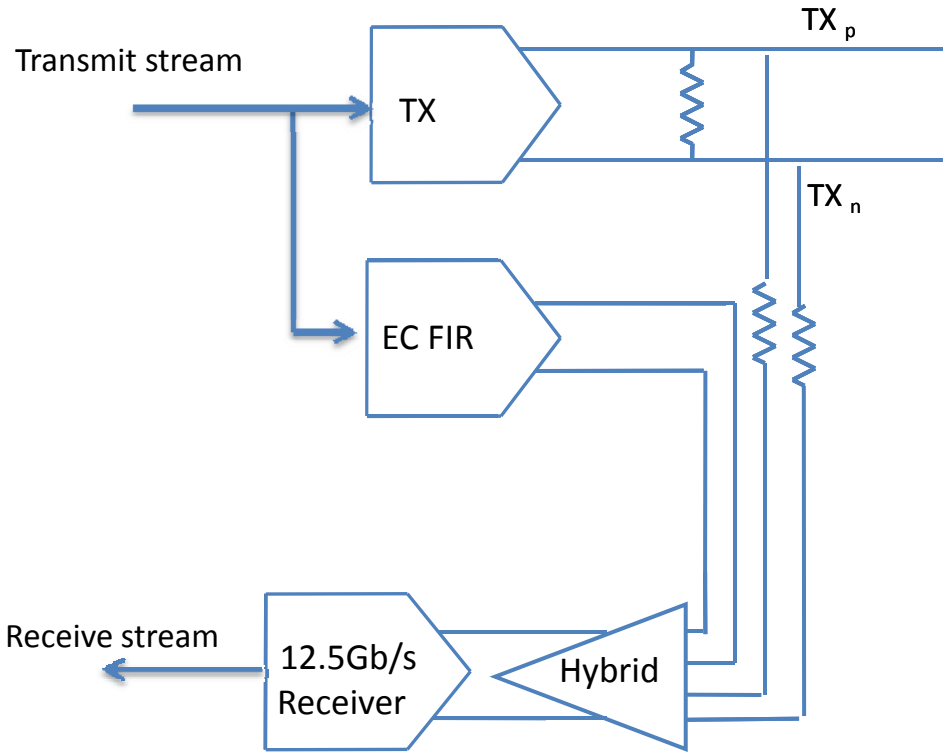

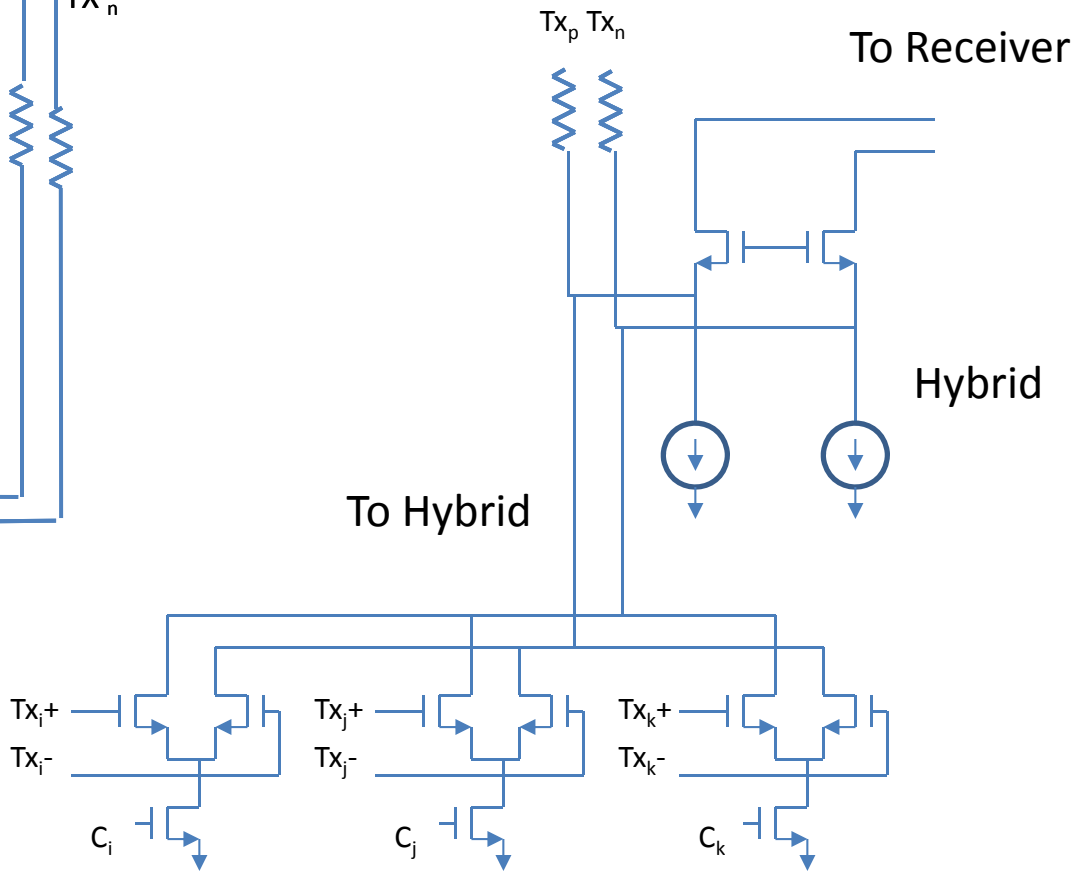

- Circuit shown is illustrative
  - More than 3 taps will be needed
- EC coefficients will be set adaptively

Key implications of full duplex 8-lane operation

- Number of transmitters and receivers is doubled relative to 4-lane half duplex operation
- Two new circuit blocks are required (hybrid and echo canceller)
- Transmitter and receiver have to be synchronous (requires master-slave operation)
- The 100G is now being distributed over 8 lanes and this may require a new gearbox mode
- Crosstalk could get impacted as each lane contributes NEXT
- Despite above factors, this is attractive because the operating SNR requirement drops substantially relative to PAM4

Complexity

- Since each lane will operate at ~12.5Gb/s with PAM2, the complexity of the transmitter and receiver (excluding the hybrid and the echo cancelling FIR filter) can be easily extrapolated from that of a 10GBASE-KR lane
- As for the hybrid and the echo canceller
  - Echo must be reduce to ~40db below TX signal; achieving required linearity in the hybrid circuit is feasible
  - Length of FIR required depends on return loss (echo impulse response) of channel and is under investigation
  - Leakage from EC to TX path will have to be minimized (>~30dB isolation) to prevent this path from adding to the tx jitter

Echo observations

- Echo response typically shows significant peaks that should be associated with channel discontinuities – should look like a TDR plot
  - Conversion from S-Parameters to time domain can sometimes generate artifacts that are non-causal
- Most channels seen have a high pass structure to the echo response
- We assume excitation by a limited bandwidth pulse and some low pass filtering in the receive path
- The taps are assumed to be non-contiguous

Simulation results

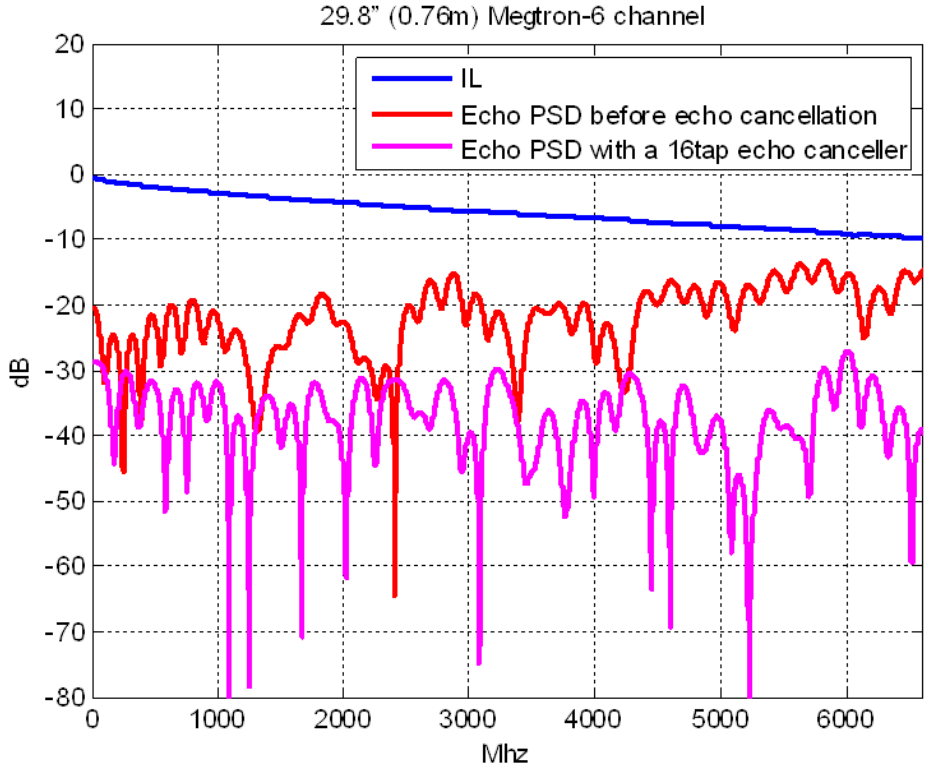

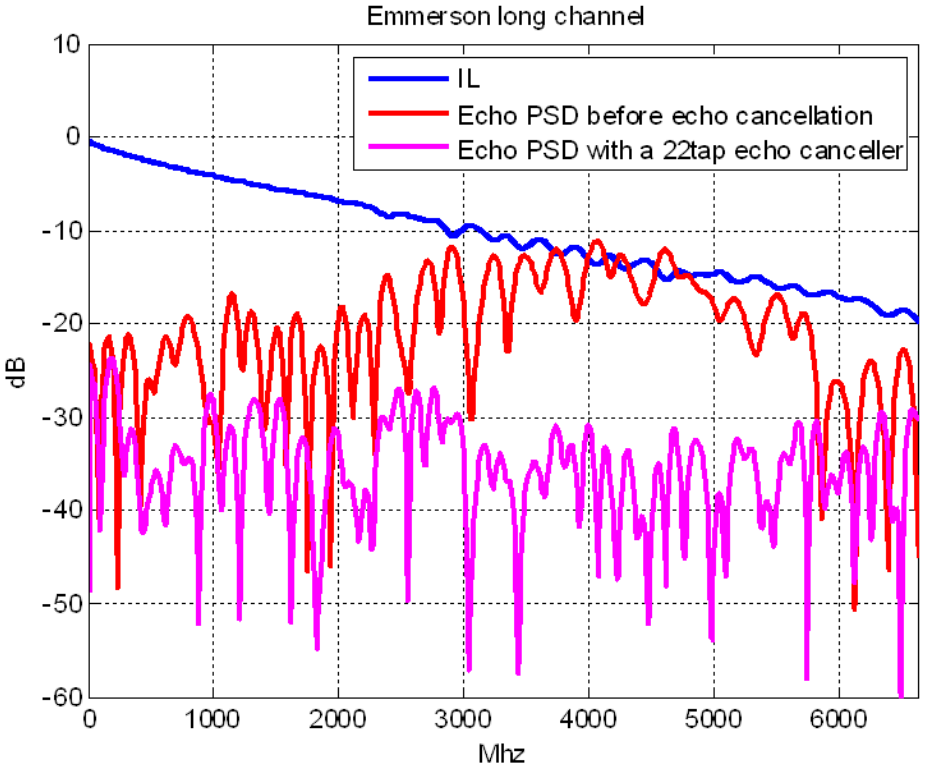

- Shows residual echo after cancellation
  - Does not address adaptation
- Indicates that a relatively small number of non-zero EC taps will provide adequate cancellation
- Other items that require study
  - Impact of loop timing on jitter
  - Impact of leakage from EC path on TX jitter

Backup

Extract the time domain oversampled S11 response from the data.

Apply low-pass filter to the echo signal to reduce out of band noise contributions.

Down-sample the S11 to Fs = 6.625Ghz

Simplistic procedure used: eliminate the N largest taps of the down-sampled S11. Take the residual taps of the echo and plot the spectrum.
Salz SNR will provide rough estimate of link margin