Enabling Improved DSP Based Receivers for 100G Backplane

Dariush Dabiri

802.3bj Task Force

IEEE 802.3 Interim September 2011

Agenda

- Goals
- Introduction
- Partial Response Channel (PRC) Signaling
- Quasi-catastrophic Trellis
- Precoding
- MLSD vs Symbol-by-Symbol Detection
- Critical Path of the Viterbi Detector
- Double Terminated Trellis
- Simulation Results
- Conclusion

Goal

- It has been shown that the DSP based PAM4 solution can cover the vast majority of the installed base backplanes with 6dB SNR margin (parthasarathy_01a_0511.pdf).
- Our goal is to show that we can increase the link margin to about 9 dB with a minor modification at the transmitter.
- Increase in SNR margin will allow more reliable operation over worst case channel and better power/performance tradeoffs.

Introduction

- For M-PAM signaling, Maximum Likelihood Sequence Detection (MLSD) will provide 3 dB gain over symbol-by-symbol detection, if M is a small number .
- Both MLSD and symbol-by-symbol detection (e.g. DFE) suffer from error propagation for worst case BP channels.
- A simple, low cost 1/(1+D) pre-coding:
  - mitigates error propagation from the DFE based solution.
  - eliminates ‘quasi-catastrophic’ paths from the MLSD trellis.
- We will show that a small modification in the transmitter allows a low power implementation of the MLSD.

Basic Partial Response Channel

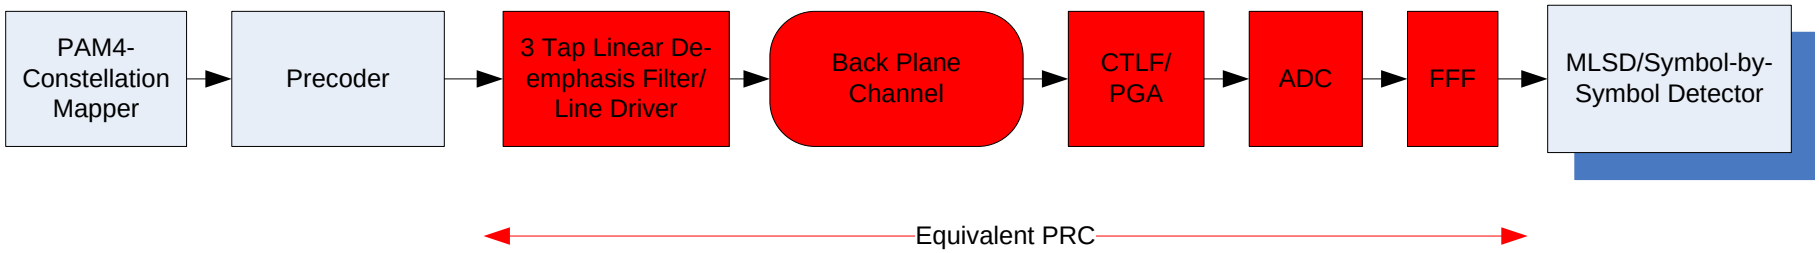

- Combination of the de-emphasis filter, the backplane channel, the CTLF and the FFF (Feed Forward Filter) creates an equivalent channel which approximates a simple 1+D channel.
- At the output of the FFF one can employ either:
  - Symbol-by-Symbol Detector:
    - Implemented by a DFE with small tap weights.
  - MLSD using Viterbi Detection
  - A combination of both

MLSD Trellis

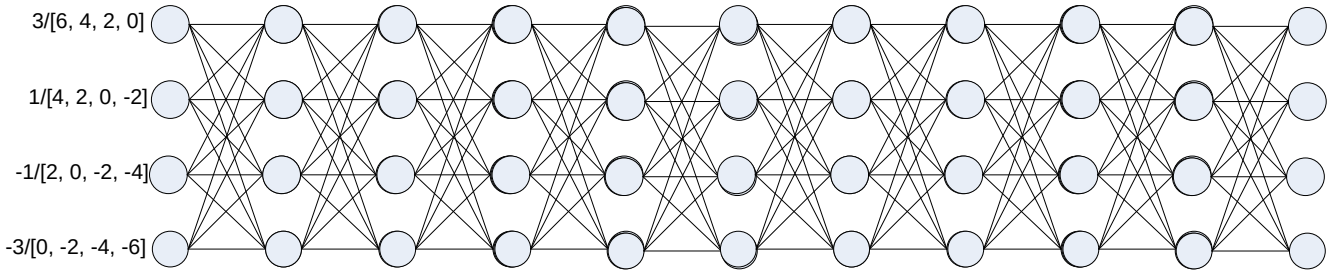

Trellis Structure for Precoded MLSD

- Each state signifies the previous input to the channel.
- Each branch signifies the output of the channel.
- Each section of the trellis has 4 states and 16 branches.

Quasi-Catastrophic Trellis

- ‘A trellis is quasi-catastrophic if it contains distinct states for which one or more of the output sequences that start from those states are identical for all subsequent time, but for which the total probability of all such non-distinguishable sequences is zero’, Forney and Calderbank: ‘Coset Codes for Partial Response Channels: or, Coset Codes with Spectral Nulls’, IEEE Trans. IT, Nov 1989.
- The trellis of the (1 + D) channel is ‘quasi-catastrophic’: all-zero output sequence can start from any state.
- For a precoded channel, the output sequence is congruent to the input of the precoder modulo M.

MLSD vs Symbol-by-Symbol Detection

- It is well known that MLSD offers approximately 3 dB extra gain compare to the symbol-by-symbol detection for the 1+D channel.
- If the impulse response of the ‘equivalent channel’ deviates from 1 + D, the MLSD gain will be less, but never worse than symbol-by-symbol detection.
- The MLSD requires implementing the Viterbi algorithm or one of its approximations.
- The Viterbi algorithm has the critical loop of ‘Add-Compare-Select’ (ACS) which makes it hard to implement in high speed applications.

Viterbi Algorithm

- An example:

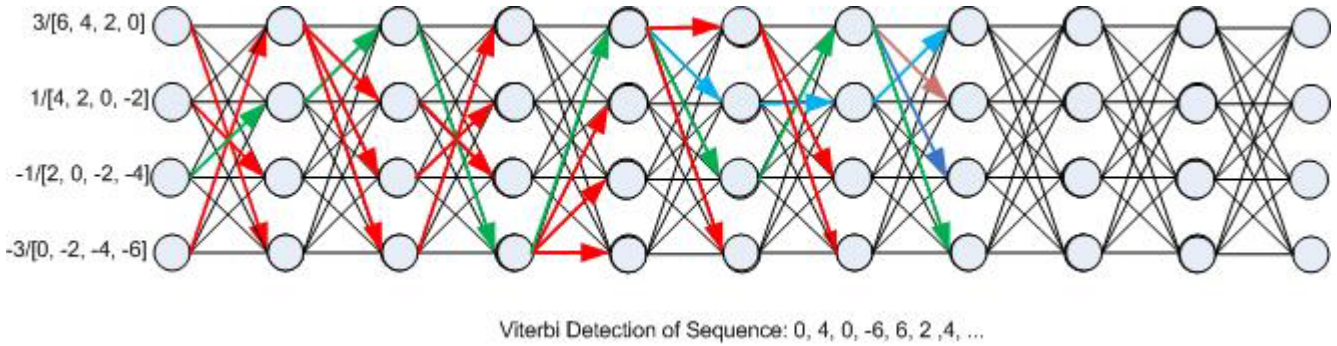

Viterbi Detection of Sequence: 0, 4, 0, -6, 6, 2 ,4, ...

- One needs to recursively computes metrics for each state at each stage of the trellis:
  - metric(s) = min(metric(s') + metric(s → s'))
  - Only the path s'→s which takes on the minimum value survives.
- The logic implementation of the algorithm at 14G is challenging: linear speed up with number of processing units.

Block MLSD: Double Terminated Trellis

- In double terminated trellis the two ends of the trellis are known at the receiver:

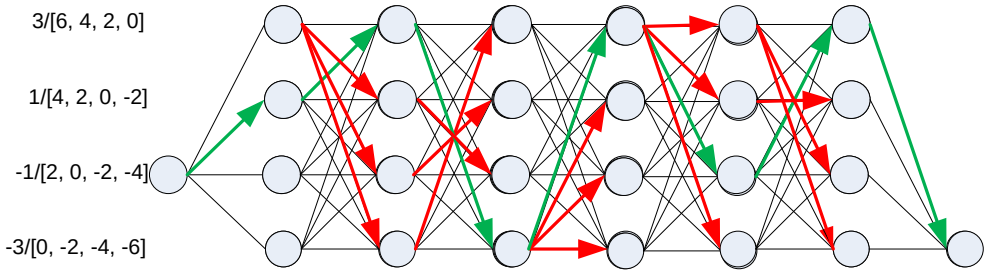

Viterbi Detection of Sequence: 0, 4, 0, -6, 6, 2 ,0

- Note that the complete detection can be achieved by the information available within the block boundary defined by termination symbols.
- Therefore one can allow implementation of the detection of multiple blocks to overlap in time.

Complete Model

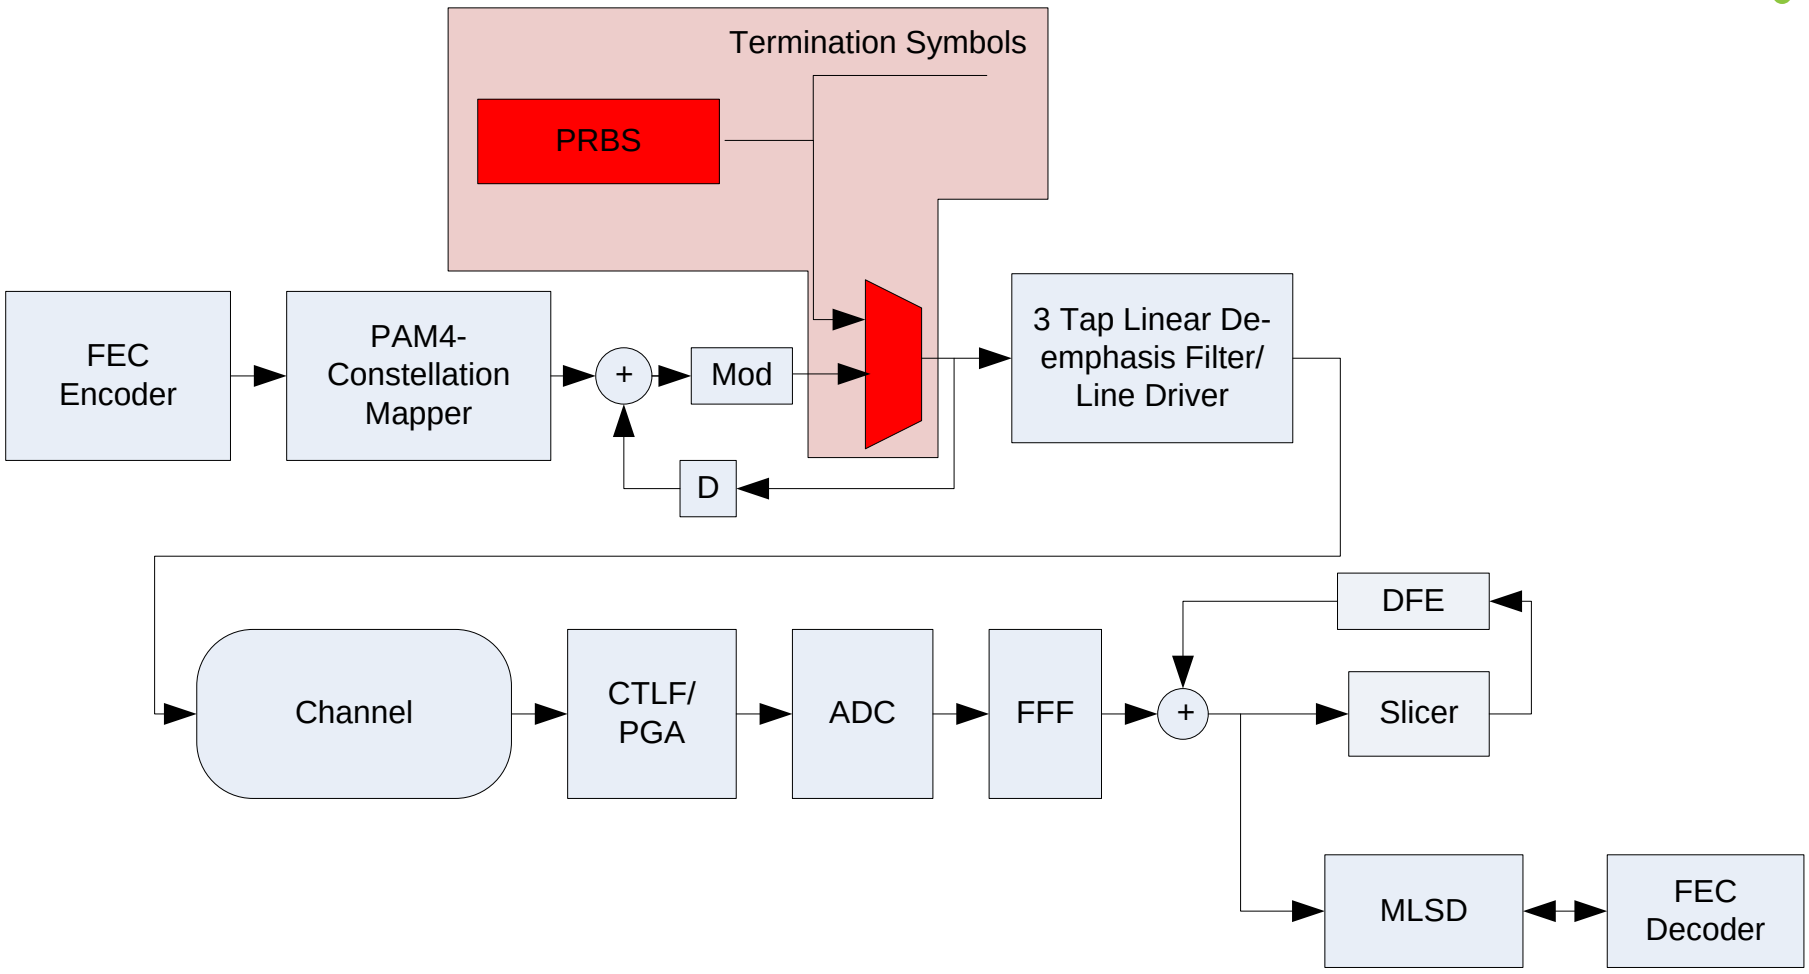

Termination Symbol

- The PRBS sequence is locked between the transmitter and receiver during the start-up:
  - It is known to both ends.
- The termination channel symbol corresponds to exactly one PAM4 symbol before precoder:
  - At the transmitter, one can add the termination symbol after the precoder, and remove it after the detector.
  - End-to-end frame structure will not be affected by insertion and deletion of the termination symbol.
- One possible choice is to add one termination symbols for N = 255 symbols:
  - Corresponds to 0.39% over-clocking
    - SNR loss of around 0.05 dB due to over-clocking.

Simulation Results Before FEC Decoder

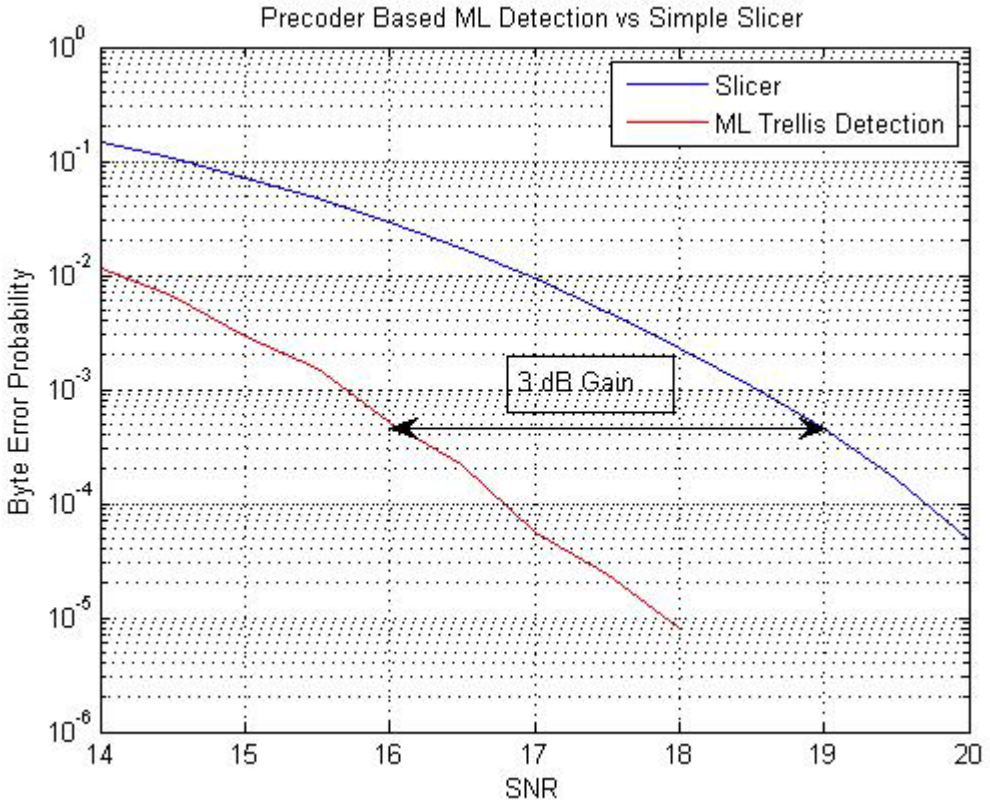

Complexity

- Block processing allows an straight forward implementation of a 4-state Viterbi detector to work at 100G:
  - Alternatively one has to unroll the ACS loop:
    - Fettweis and Meyr, 'High-rate Viterbi processor: a systolic array solution', IEEE JSAC, Oct, 1990.
- Synthesis results indicates that the power/area of the detector is less or comparable with the FEC decoder
  - Based on a 28nm node library.

Conclusion

- We can increase the performance of the DSP based implementations by about 3 dB with a very minor modification in the transmitter and relatively small increase in power at the receiver.
- The proposed change can easily be adapted to all proposed FEC frame structures.

Thank you!