Channel Code Considerations for 10GbT Signaling

January 2004
Jose Tellado, Dariush Dabiri, Teranetics
Ofir Shalvi, Ariel Yagil, TI

Supporter: Scott Powell, Broadcom

Motivation

- Two different signaling architectures have been discussed so far
  - – TCM: 10PAM 4D-8State code at 833Ms/s – Solarflare
  - – LDPC: 8PAM (1723,2048) code at 1000Ms/s – Intel
- These proposals differ in other aspects too
    - Packet format, overhead bits etc.
    - Equalization approach
- This presentation compares the coding schemes while normalizing other factors out
- We also include some other well known schemes

Teranetics

TEXAS INSTRUMENTS

Ideal Performance Bounds

- A “good” transceiver design would convert the ISI+Xtalk+noise channel into an (approx) AWGN channel
- Shannon capacity for and ideal AWGN channel
  - $C = ½*\log_2(1+SNR)$ bits per 1D symbol (bps)
- For example, a capacity approaching code (“infinite” delay) can operate at 2.5bps with “zero” BER at SNR of 15 dB

Practical Performance Bounds

- For uncoded MPAM, M even
  - BER ~ Q(sqrt(3/(M^2-1)*SNR))
- Solving the equation above we have
  - Rate = ½*$\log_2$(1+SNR/G(BER)) per 1D symbol
  - G(BER) is the Gap or Loss relative to capacity and depends on the target BER.
- For uncoded PAM, G(10^{-12})~12dB
- For coded systems, d increases and G(BER) is reduced
  - The reduction of G(BER) is called coding gain

Multi-channel SNR in code design

- Note that the IL and residual X-talk (NEXT, FEXT, ANEXT) level may vary from pair to pair due to the physical channel parameters (cable length, separation, connectors, …)
- SNR variations should be considered in the code design

Possible Channel Codes

Scheme	Loss from Capacity @ BER=1e-12	Delay [µsec]	Maturity	Receive Complexity	Equalizer
TCM + shaping	5-8dB (Solarflare TCM=8dB)	< 0.1	Mature	Low – moderate	Pre-coding or receiver equalization
TCM+RS +Shaping	3-6dB	0.75-4.5	Mature	Low – moderate	Pre-coding recommended
LDPC + shaping	1-4dB (Intel LDPC=3.8dB)	0.5-2	New	High	Pre-coding

- Lower loss from capacity translates to better link quality
  - For each scheme, lower loss typically requires a more complex receiver and more delay (latency)

The 1GbT 5PAM-4D-TCM code

Achieves BER=1e-12 for 2 bps in an ideal (no ISI) AWGN channel at Es/N0 (dB) of

$$23.9 - (5.7 - 10\log_{10}(8/5)) - 0.4 = 19.9 \text{ dB}$$

SNR required by un-coded 4PAM

Coding gain = 3.66 dB

(power ratio between 5PAM and 4PAM)

Shaping gain

Gap from capacity = 8dB

LDPC/Turbo Codes

- A large body of work (most starting mid 90s) has shown that LDPC/Turbo codes can approach the Shannon bound
  - Most of the published literature has focused on the low SNR
- Intel's LDPC 8PAM (1753,2048) proposal:
  - Achieves BER=10^{-12} for 2774/1024=2.68 bps and SNR=19.8dB
  - The Shannon bound at 2.68 bps is SNR=($2^{2*2.68}$-1) ~ 16dB
  - The loss from capacity is 19.8-16 = **3.8dB**
  - The intrinsic decoding delay (i.e. with infinite HW) is 0.256 micro-seconds, but practical decoders will have additional delay
  - The SNR margin could be improved by using shaping algorithms, different 4D mappings, larger block sizes, etc. typically at the expense of more complexity and/or latency

Performance of current Proposals

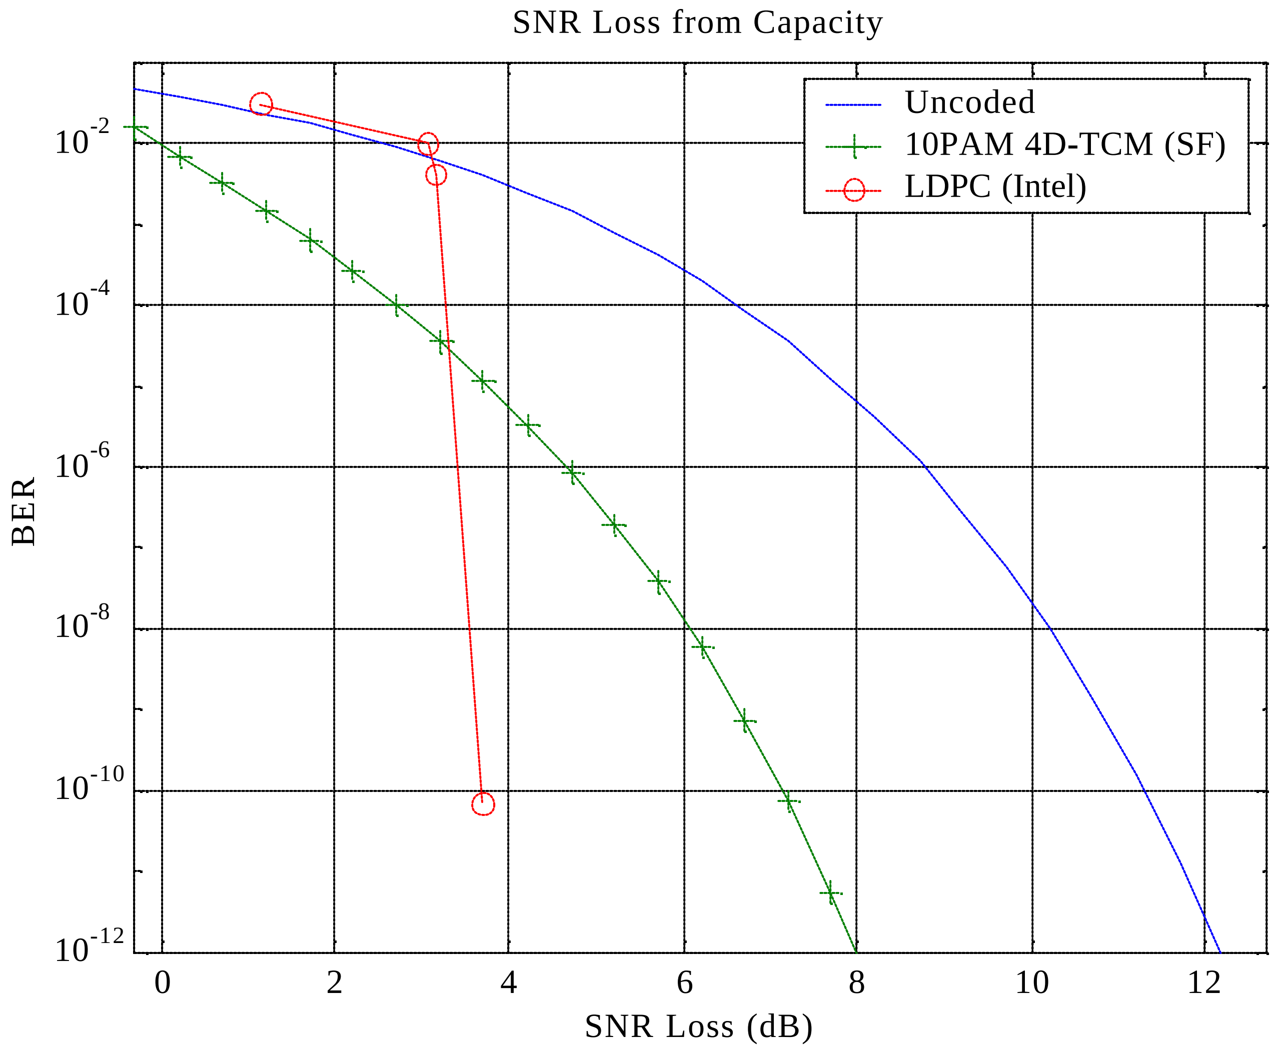

A concatenated RS+TCM based on the 1GbT code

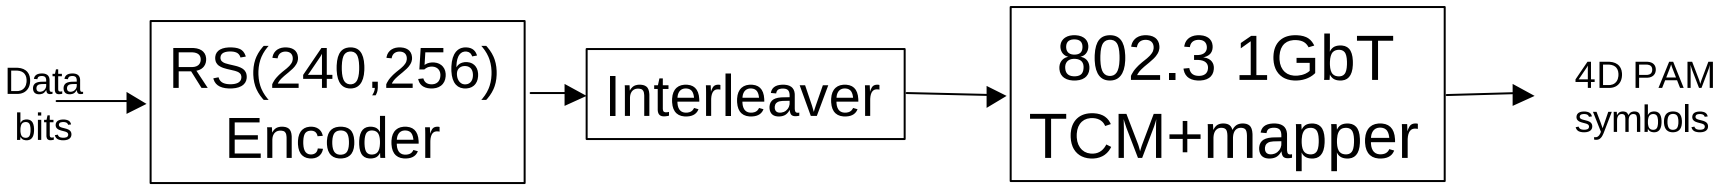

- Data rate of 1.875 bits per dimension
- Can easily be generalized to 10PAM (2.8125 bits per 1D-PAM symbol)
- Low complexity, mature decoding algorithms

Teranetics

Texas Instruments

Concatenated RS+TCM – Performance Analysis

- BER=1e-12 @ SNR of 15.4 dB – **only 4.4 dB short of Shannon Capacity**
- Using a standard hard decoding algorithm
- Similar gain to that of the LDPC proposed in the November meeting
- Analysis assumes ideal interleaver

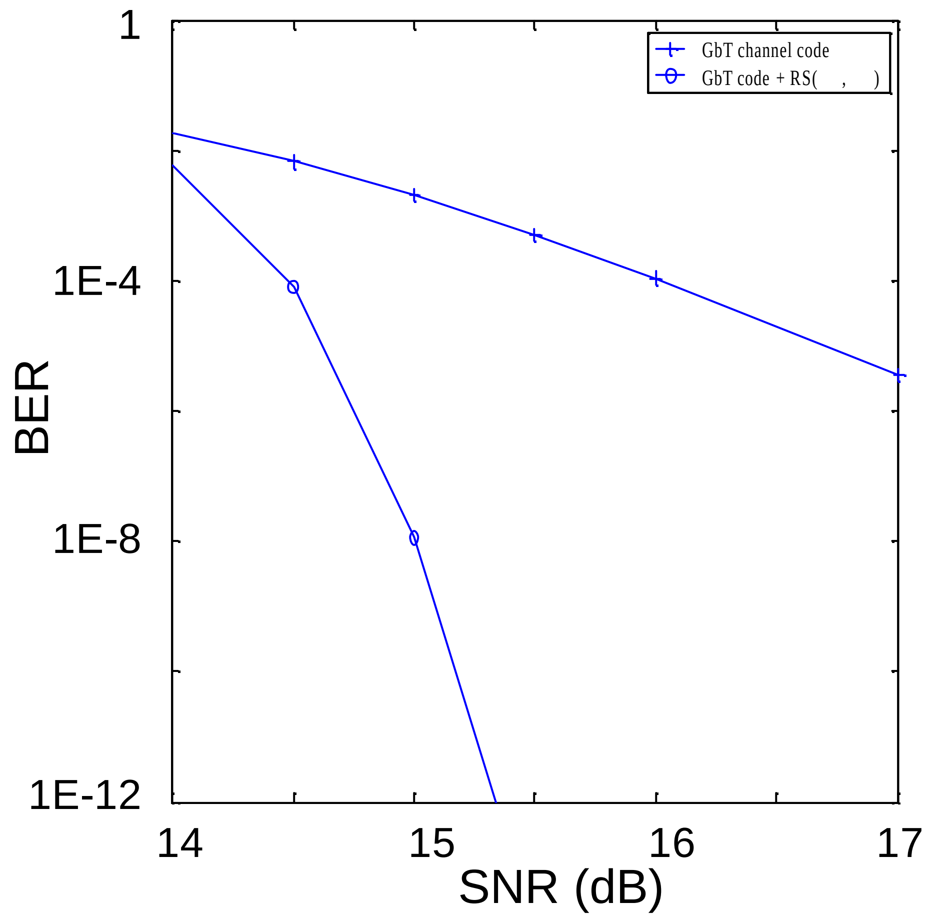

Concatenated RS+TCM – Performance/Latency tradeoff

- The analysis of the code with a short interleaver is more complicated.
- We assess that the loss due to an interleaver latency of 2.5 micro-seconds is a small fraction of a dB.
- We assess that the loss due to a latency of 1 micro-second is about 1-1.5 dB.

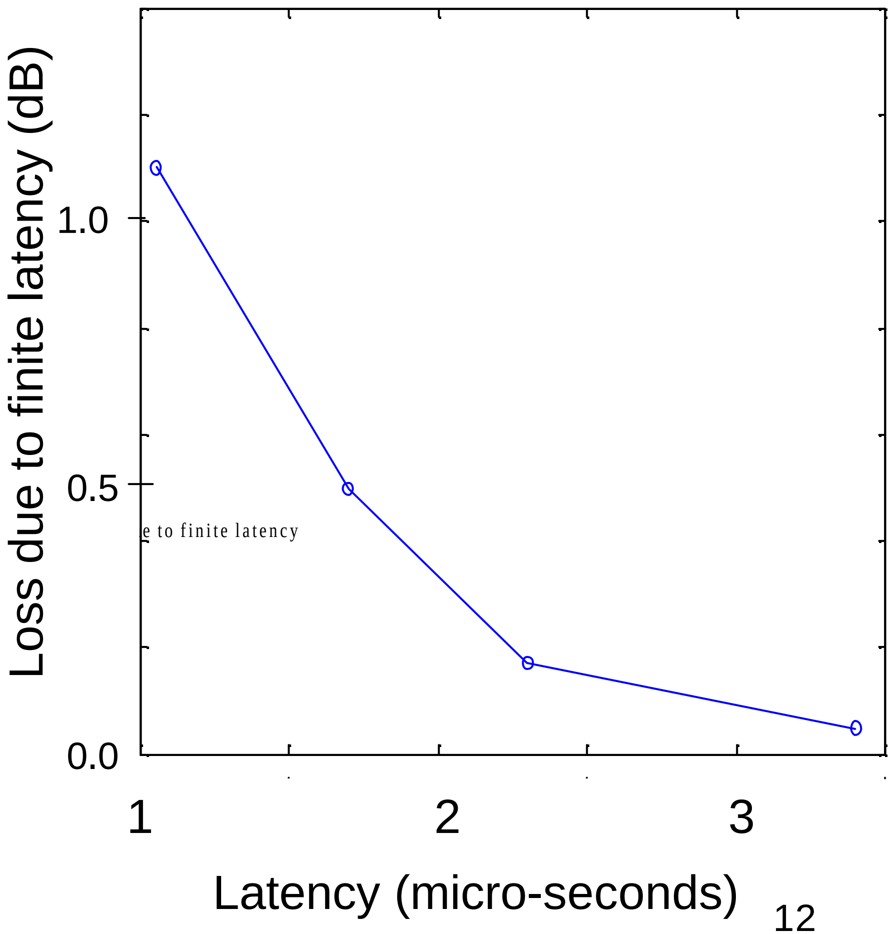

Improved concatenated codes

- It is possible to further gain 0.5-0.7 dB by employing constellation-shaping algorithms.
- Lower latency or additional coding gain can be achieved by employing other concatenated coding schemes.

Performance-Complexity Summary

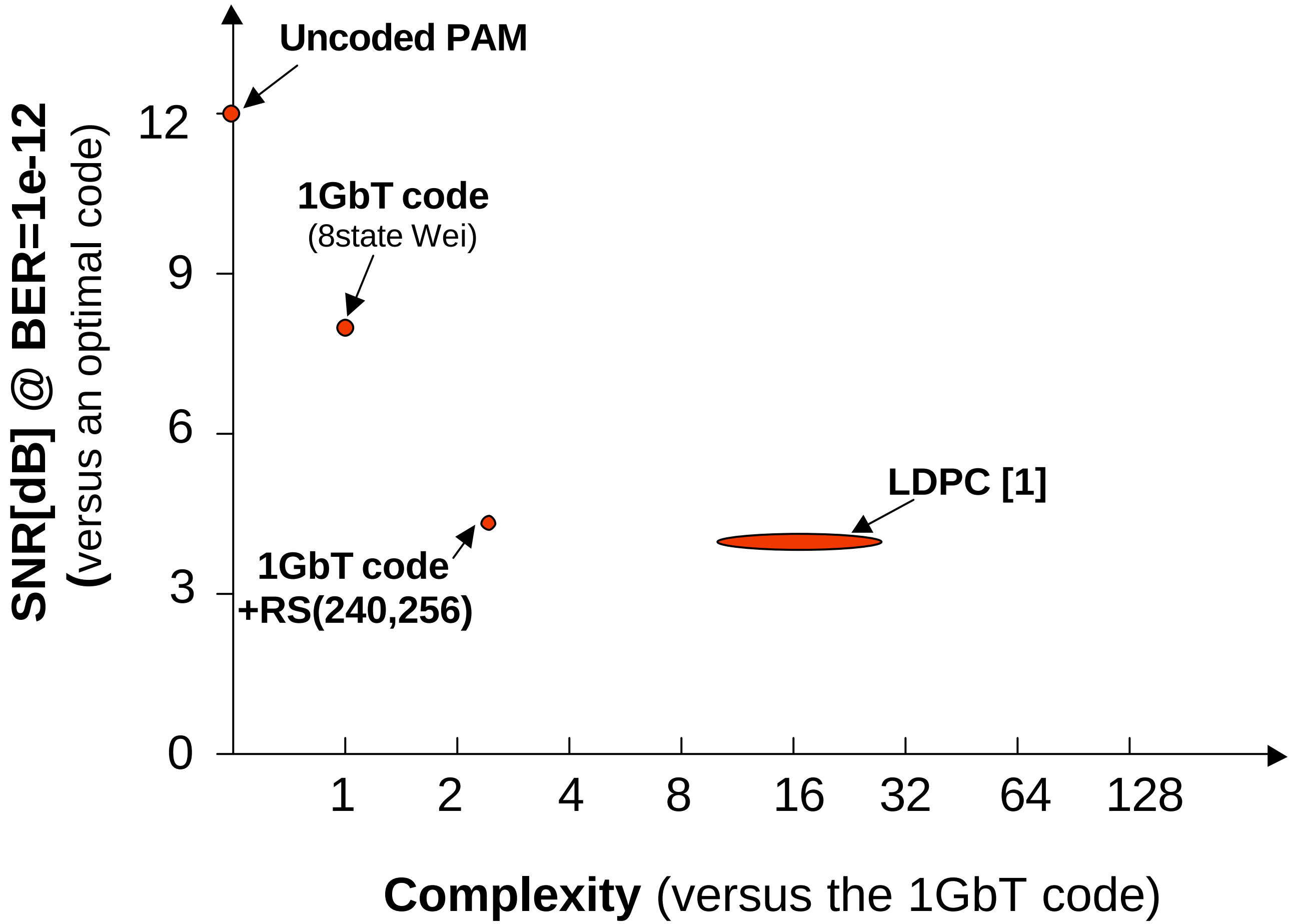

Performance-Complexity Summary

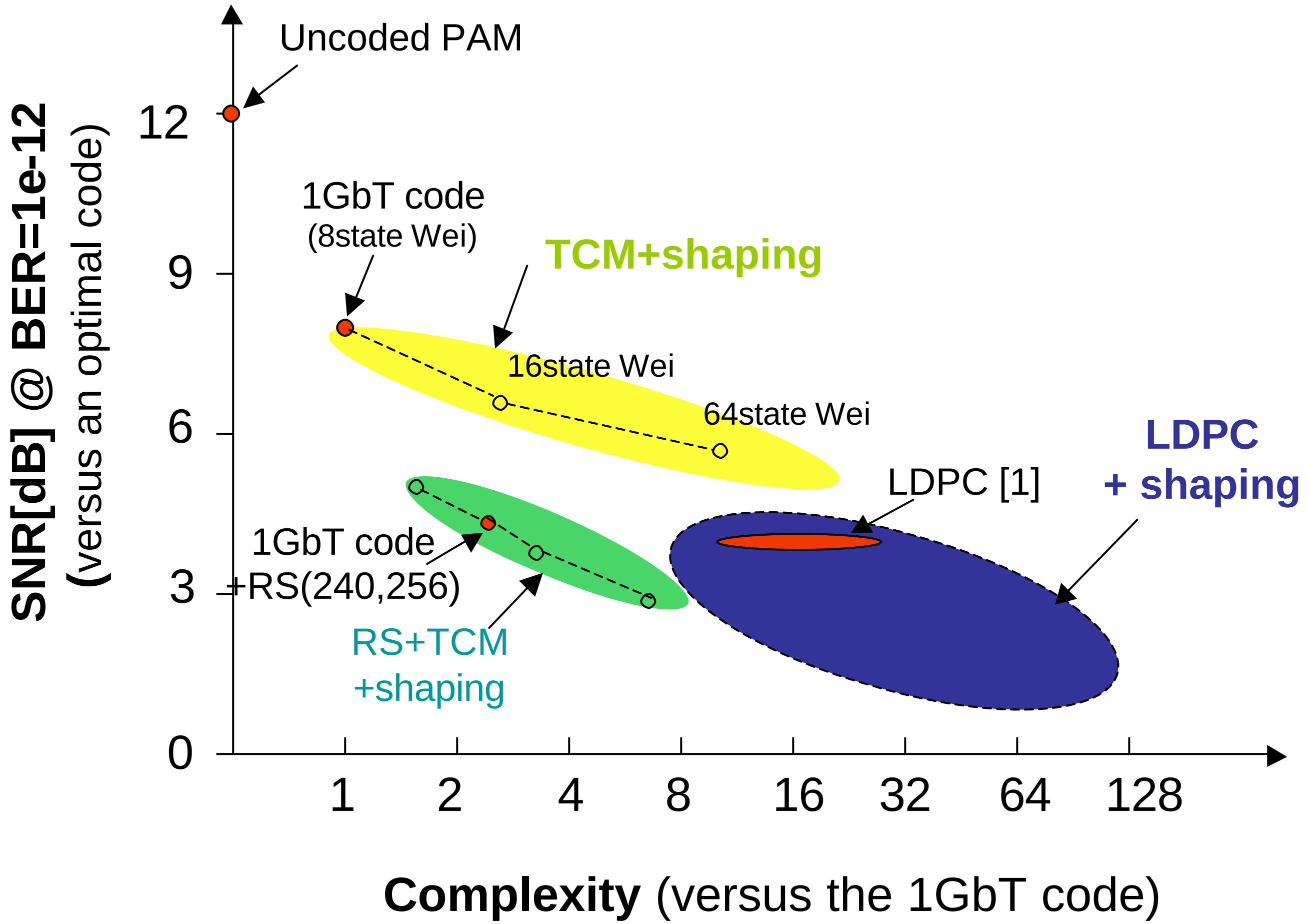

Conclusion

- We seek input from task force participants on:
  - Latency budgets
  - Performance/complexity tradeoffs
- Based on these inputs, specific codes can be optimized for the 10GBASE-T application
  - Concatenated Codes
    - Optimize for a tolerable latency range
  - LDPC
    - Optimization and more detailed evaluation of performance and complexity.
  - Optimize symbol rate and packet overhead
  - Evaluate addition of constellation-shaping gain to codes