



Forward Error Correction (FEC) for EPoC

Christian Pietsch – Qualcomm
David Barr – Entropic

Supported by

- Stefan Brueck (Qualcomm)
- Jorge Salinger (Comcast)
- Rick Li (Cortina)
- Lup Ng (Cortina)


EPoC Objectives – A Motivation for FEC

- As part of the agreed EPoC objectives for 802.3bn, there is:
 - *“Provide a physical layer specification that is capable of a baseline data rate of 1 Gb/s at the MAC/PLS service interface when transmitting in 120 MHz, or less, of assigned spectrum under defined baseline plant conditions”*
 - *“PHY to have:*
 - *a downstream frame error rate better than 10^{-6} at the MAC/PLS service interface*
 - *an upstream frame error rate better than 5×10^{-5} at the MAC/PLS service interface”*
- In order to be able to meet very low error rate targets, while providing high spectral efficiency, Forward Error Correction (FEC) techniques have been successfully used in deployed digital communication systems
- This presentation briefly reviews the benefits of FEC coding, showing how modern Low Density Parity Check (LDPC) codes can achieve high data rates very close to the Shannon limit, and proposes to adopt LDPC codes as baseline FEC in EPoC for both the downstream and upstream

Forward Error Correction (FEC) Benefits


- At the transmitter, FEC adds redundancy to the information in a controlled way in order to detect and correct transmission errors at the receiver
 - This approach allows achieving very high reliability in data transmission
- FEC is well suited for EPoC to achieve the target error rates at the MAC/PLS interface at the given SNR
- State-of-the-art coding schemes, like LDPC codes, have been successfully used in commercial systems
- Instead uncoded transmission would require significantly higher SNR to achieve the same spectral efficiencies (~ 6 dB more), thus, increasing the transmit power and power consumption significantly

FEC Comparison



Coding-Gain: LDPC vs. Reed-Solomon

LDPC (4608,3648) versus RS(255,203)x2 for 256QAM



Discussion

- The previous slides compare LDPC codes, Reed Solomon codes and uncoded transmission in terms of SNR [dB] and spectral efficiency [bits/s/Hz] in AWGN conditions
 - The BER (after decoding in case of FEC) is fixed to $1e-6$
 - The theoretical upper limit (Shannon curve) is given as well
- Results
 - For a given spectral efficiency, the LDPC code outperforms the Reed Solomon code by roughly 3-4 dB
 - For a given spectral efficiency, the LDPC code outperforms uncoded transmission by roughly 6 dB (~ 3dB better than the Reed Solomon code)
 - For a given SNR, the LDPC code improves the spectral efficiency compared to uncoded transmission by roughly 2 bits/s/Hz
- Conclusion
 - State-of-the-art FEC schemes allow improving spectral efficiency and/or power efficiency in EPoC considerably, fulfilling EPoC BER requirements

Baseline Proposal #1

- The PHY shall implement and use a mandatory Forward Error Correction (FEC) scheme for both downstream and upstream transmissions

- Moved by:
- Seconded by:
- Technical motion ($\geq 75\%$)
- Yes / No / Abstain

LDPC as State-of-the-Art FEC

■ LDPC \equiv Low-Density Parity Check

- Invented in 1960, but impracticably complex for the time, then forgotten
- Invention rediscovered in 1996
- Now extensively studied in academic literature
- Well-known to have excellent performance
 - » approaches Shannon channel capacity limit
 - » and very low error floor (to be explained)
 - can achieve very low block error rates (BLER)
 - without necessarily relying on retransmissions or outer codes
- LDPC became FEC of choice for new communication standards
 - » Wi-Fi (802.11n, .11ac, .11ad, WiGig) and WiMAX (802.16e)
 - » Ethernet 10GBase-T (802.3an) and DVB-S2

Quasi-Cyclic LDPC Block Codes

- LDPC Block Codes are defined by a sparse Parity-Check Matrix
 - Sparse means matrix is primarily populated with zeroes
 - These code matrices are essentially discovered, e.g., via computerized search
 - » then extensive computer testing needed to prove good performance
 - Performance can be tested against AWGN (additive white Gaussian noise)
 - » yielding the basic block error rates (BLER) vs. SNR (see next slide)
 - Performance can be tested against non-Gaussian noises (e.g., interferers)
 - » Entropic did extensive robustness testing with simulated interference sources
- Encoding is trivial
 - e.g., CNUs can implement a low-complexity upstream FEC
- Decoding is highly parallelizable in hardware
 - Easily partitioned among multiple parallel HW execution units
 - much more so than alternative Turbo Codes
 - Supports >Gbps throughput, while keeping HW clock-rates low
 - Fully exposed to Moore's Law cost-reductions (implementations entirely digital logic)
 - Decoding algorithms continue to improve (e.g., belief propagation)

Codeword Length

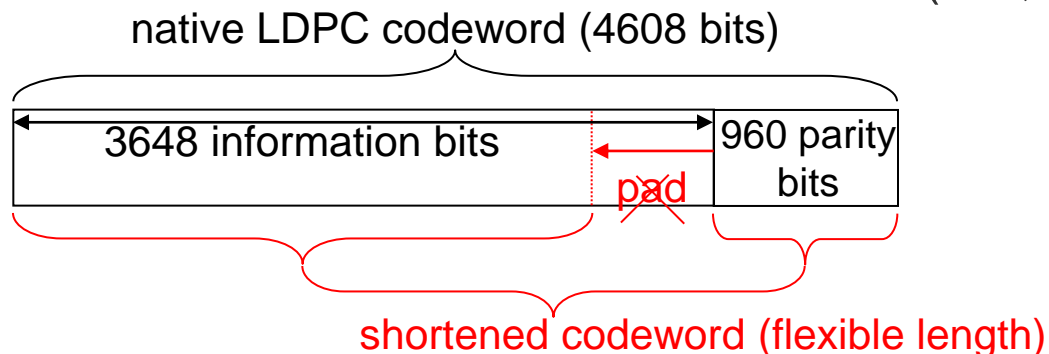
- The usage of Long and Small codewords was suggested (boyd_01a_0712.pdf)
 - e.g., Long (~4.5k bits) codewords for the downstream (~90% code-rate)
 - e.g., Smaller codewords for the upstream (~80% code-rate)
 - » with codewords shortened for additional efficiency
- Entropic proposes
 - Long Code ~4.6k bits at 85% code-rate
 - Small Code length ~1/8th the long code
 - Parallelism is 1/8th the long code
 - Enables same HW to decode (in the same amount of time = one symbol period):
 - » one long codeword
 - » or 8 small codewords simultaneously
 - this is **very** useful for OFDMA (i.e., simultaneous CNU transmitters)
 - hardware reuse avoids implementing 8 separate decoders, and does not increase die area
 - Optimally shortened codewords
 - Optimized on a payload-by-payload basis
 - With all codewords in a payload shortened ~equally
 - » Not just the last codeword
 - » Payload is better protected, overall
 - achieves lower BLER than targeted

Realizing Various Codeword Lengths

- It is highly beneficial to have a variety of codeword lengths available
 - e.g., to adapt to particular payload sizes, traffic patterns,...
 - e.g., to adapt to channel-conditions (which may vary in time)
- Three approaches to implementing a variety of codeword lengths:
 - Design a **set** of specific LDPC mothercodes (each with its own parity-check matrix)
 - Shortening of codewords from any given LDPC mothercode (see next slide)
 - Puncturing of codewords from any given LDPC mothercode (see next+1 slide)
- IEEE 802.11n specifies a **set** of three different LDPC mothercodes
 - Each with its own parity-check matrix:
 - Long codeword with native length 1944 coded bits
 - Medium codeword with native length 1298 coded bits
 - Small codeword with native length 648 coded bits
 - These LDPC mothercodes should also be considered for 802.3bn EPoC PHY

Shortening of LDPC Codewords

- Codewords shorter than native length can be easily constructed:
 - transmitter loads some **reduced** number of information bits
 - padding bits (binary 0s) loaded to fill native info block size
 - transmitter calculates the native (fixed) number of parity bits
 - padding bits (known to the receiver) not transmitted
 - receiver inserts the untransmitted padding bits
 - receiver decodes native-length codeword
 - receiver extracts the shortened number of corrected info bits
-
- Yields a lower-than-native code-rate (i.e., higher coding-gain)



Puncturing of LDPC Codewords

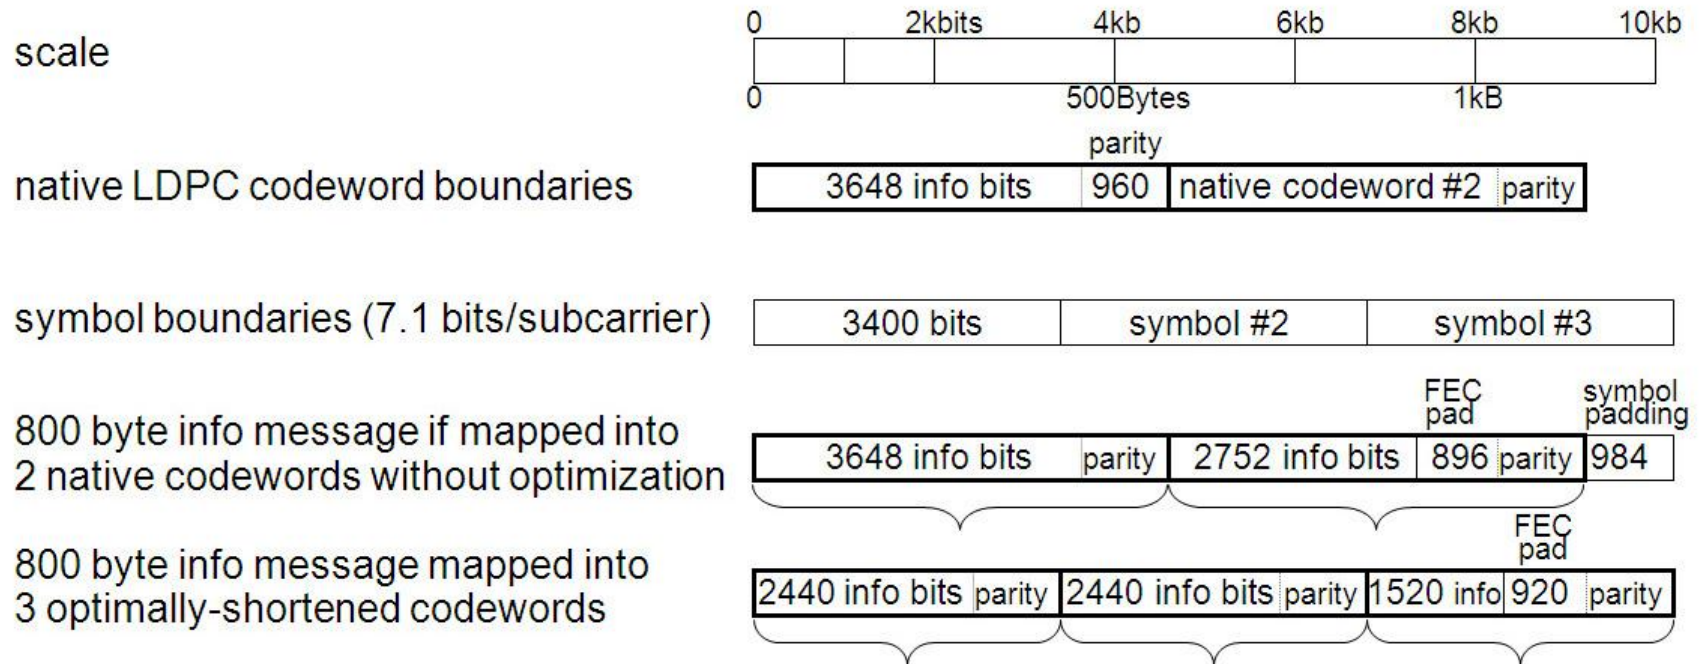
Another method of implementing adaptive FEC coding rates

- e.g., in response to time-varying channel state information
 - also reuses a single LDPC mothercode
-
- Start with normal encoding of LDPC codewords
 - Transmitter [optionally] does not send some of the parity bits
 - i.e., the parity field of the codeword is *punctured*
 - ~equivalent to a codeword with a smaller parity field
 - Receiver uses a normal LDPC decoder
 - but knows the puncturing pattern where parity bits are missing
 - Yields a higher-than-native code-rate (i.e., lower coding gain)
 - than the native LDPC mothercode

Optional Codeword Optimization per Payload

- EPoC reuses EPON MAC: i.e., MPCP GATE grant specifies exactly:
 - » # Bytes to transmit upstream
- Thus, CNU and CLT can agree on codeword optimization
 - Grant size can determine which codeword and shortening to use
 - » e.g., big grants can use the long code; small grants use the small code
- Codeword Optimization:
 - Tx & Rx agree on extent of shortening:
 - no need to manage from headend management entity
 - » to maximize number of codewords (and parity) per payload
 - without increasing the number of symbols used
 - no impact on throughput
 - » spreads info bits more equally among all codewords
 - coding gain of lower code-rate distributed over all symbols
 - increases robustness of entire FEC-protected payload
 - » occupies otherwise wasted padding with parity bits
- All codewords in given payload have same (shortened) length
 - except the last codeword (which may be shortened further)
- Reduced power consumption
 - LDPC decoder requires fewer iterations

Example of Codeword Optimization per Payload




Specific Results of Optimization:

- 3 shortened codewords (instead of 2 native-sized codewords)
 - fitting precisely within the same 3 OFDM symbols
 - symbol padding repurposed into new parity bits
- effective code-rate reduced to 0.72 from native 0.79
 - with no impact on throughput or decoding performance
 - entire payload more robust (0.7dB additional margin)
 - » at no cost!

LDPC Power Consumption Reduction

- Adaptively optimized codeword-shortening yields additional SNR margin
- Additional SNR margin facilitates LDPC decoding
 - LDPC decoder might be $\sim 1/3^{\text{rd}}$ of the PHY implementation
 - Fewer decoding iterations would be required
- Reduces power consumption
 - for example...




LDPC Code-Rate Survey

- Many standards support multiple FEC code-rates:
 - DVB-C2 0.67 ~ 0.90
 - DVB-T 0.50 ~ 0.88
 - DVB-T2 0.50 ~ 0.83
 - DVB-RCT 0.50 ~ 0.75
 - WiMAX 0.50 ~ 0.75
 - 802.11n 0.50, 0.75, 0.83 (implemented via 3 different LDPC matrices)
- Higher coding-rates yield lower coding-gains
 - Largely cancel each other out (in terms of PHY-rate)
i.e., PHY-rate vs. SNR remains ~constant
 - although higher code-rate reduces constellation order
- Trivial Reuse of LDPC Codeword Shortening
 - which is needed for last-codeword shortening anyhow
 - Provides an adjustable range of code-rates
e.g., 0.50 ~ 0.85 (i.e., up to the native rate)

Long Code Word LDPC – Example 1


- Long Codeword = 3900 info bits + 700 parity bits = 4600 bits (575Bytes)
- Native Code-Rate = 3900 bits ÷ 4600 bits \approx 85%
- Expansion Factor (HW parallelism) = 100 (# of execution units)
- Designed & extensively tested with narrowband & wideband impulse interference




Small Codeword LDPC – Example 2

- Small Codeword = 432 info bits + 144 parity bits = 576 bits (72Bytes)
- Native Code-Rate = 432 bits ÷ 576 bits = 75%
- Expansion Factor (HW parallelism) = 12 (~1/8th that of long codeword)

LDPC(576,432) with QPSK and 32QAM



BER Curves for 802.11n LDPC Code – Example 3



Summary

- Overview of LDPC as a modern FEC for EPoC
- Proposal for both Long and Small LDPC mothercodes
 - Consistent with suggestions from others
 - Enables same HW to decode (HW reuse means no increase in die area):
 - » one long codeword **or** 8 small codewords simultaneously (very useful for OFDMA)
- Proposal for Long, Medium and Small LDPC mothercodes from 802.11n
- Overview of codeword shortening and puncturing
 - Optional Codeword Optimization per Payload
 - » Repurposes otherwise wasted padding bits into parity bits
 - » Improves robustness of all codewords in the payload (not just the last one)
 - » Reduces power consumption of LDPC decoder
- Performance of proposed LDPC mothercodes shown to be excellent

Baseline Proposal #2

- The FEC shall use Low Density Parity Check (LDPC) codes for both downstream and upstream transmissions

- Moved by:
- Seconded by:
- Technical motion ($\geq 75\%$)
- Yes / No / Abstain

Baseline Proposal #3

- The code shall include both long and short FEC blocks
 - The details of LDPC code block lengths are for further study

- Moved by:
- Seconded by:
- Technical motion ($\geq 75\%$)
- Yes / No / Abstain



Thank you