



FEC Codeword Format and Alignment Mechanism

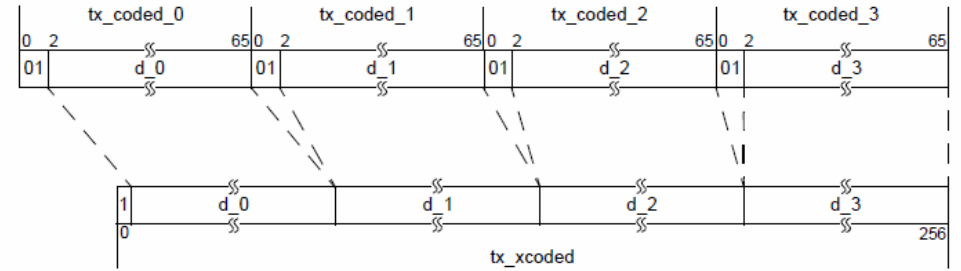
Glen Kramer

Broadcom

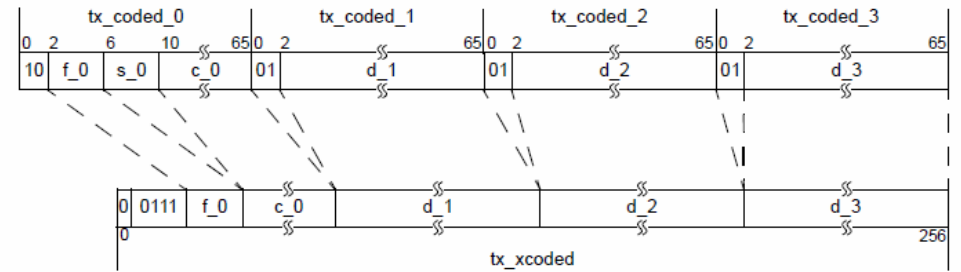


256b/257b Line Coding

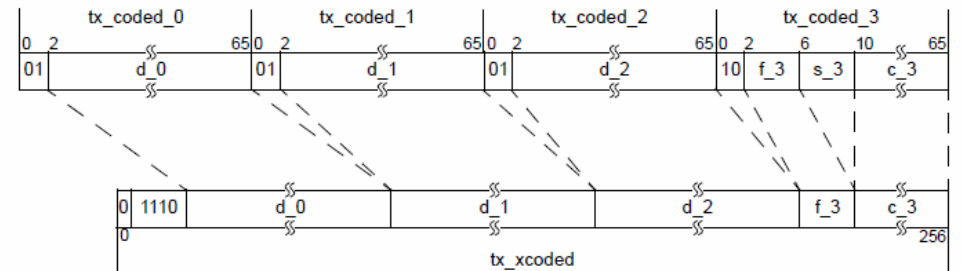
- Specified in 802.3, Clause 91 for 100GBASE-R
- Transcodes four 64b/66b blocks into one 256b/257b block
 - Throughput gain of 2.72% or 680 Mbps per lane
- The downstream LDPC Codeword accepted in Orlando (LDPC[18493,15677]) already assumes 256b/257b encoding
 - **Payload:** 61 block x 257 bits = 15677 bits
(entire 257b blocks are protected)
 - **Parity:** 11 blocks x 256 bits = 2816 bits
(only 256 bits per block are used to carry parity)
 - **Total CW size:** 15677 bits + 2816 bits = 18493 bits
or 61 + 11 = 72 blocks




Example 1: All data blocks



Example 2: Control block followed by three data blocks




Example 3: Three data blocks followed by a control block



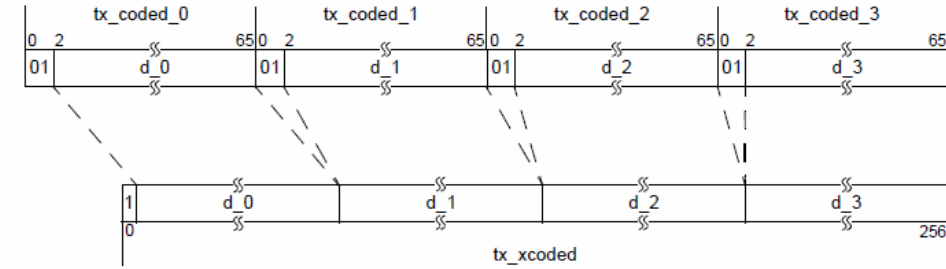
Example 4: All control blocks

Parity Delineation

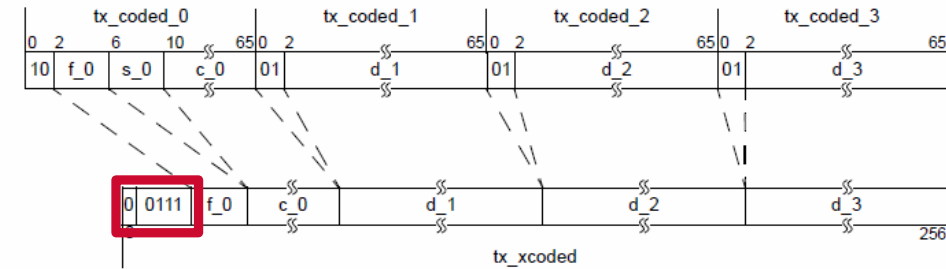
- In .3av we have defined special 64b/66b sync headers for parity
 - SH = 01: Data
 - SH = 10: Control
 - SH = 00, 11: Parity
- In 256b/257b, sync headers are collapsed to a single bit (1:All data, 0:Control+[Data])
- However, one control code-point is available to mark the parity



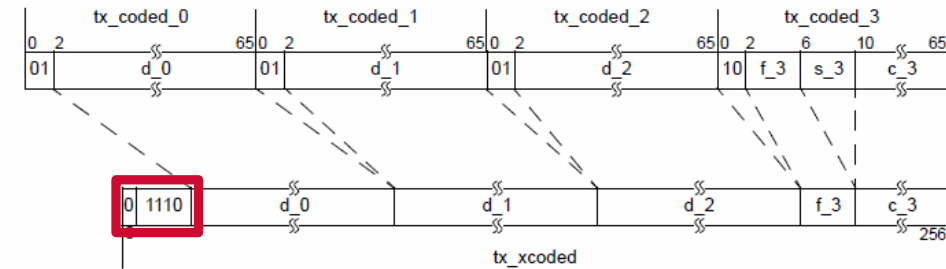
Code-point 1111_b is not used in regular control blocks.
 This code-point can mark the beginning of Parity section of FEC Codeword




Example 1: All data blocks



Example 2: Control block followed by three data blocks




Example 3: Three data blocks followed by a control block



Example 4: All control blocks

Parity Delineation

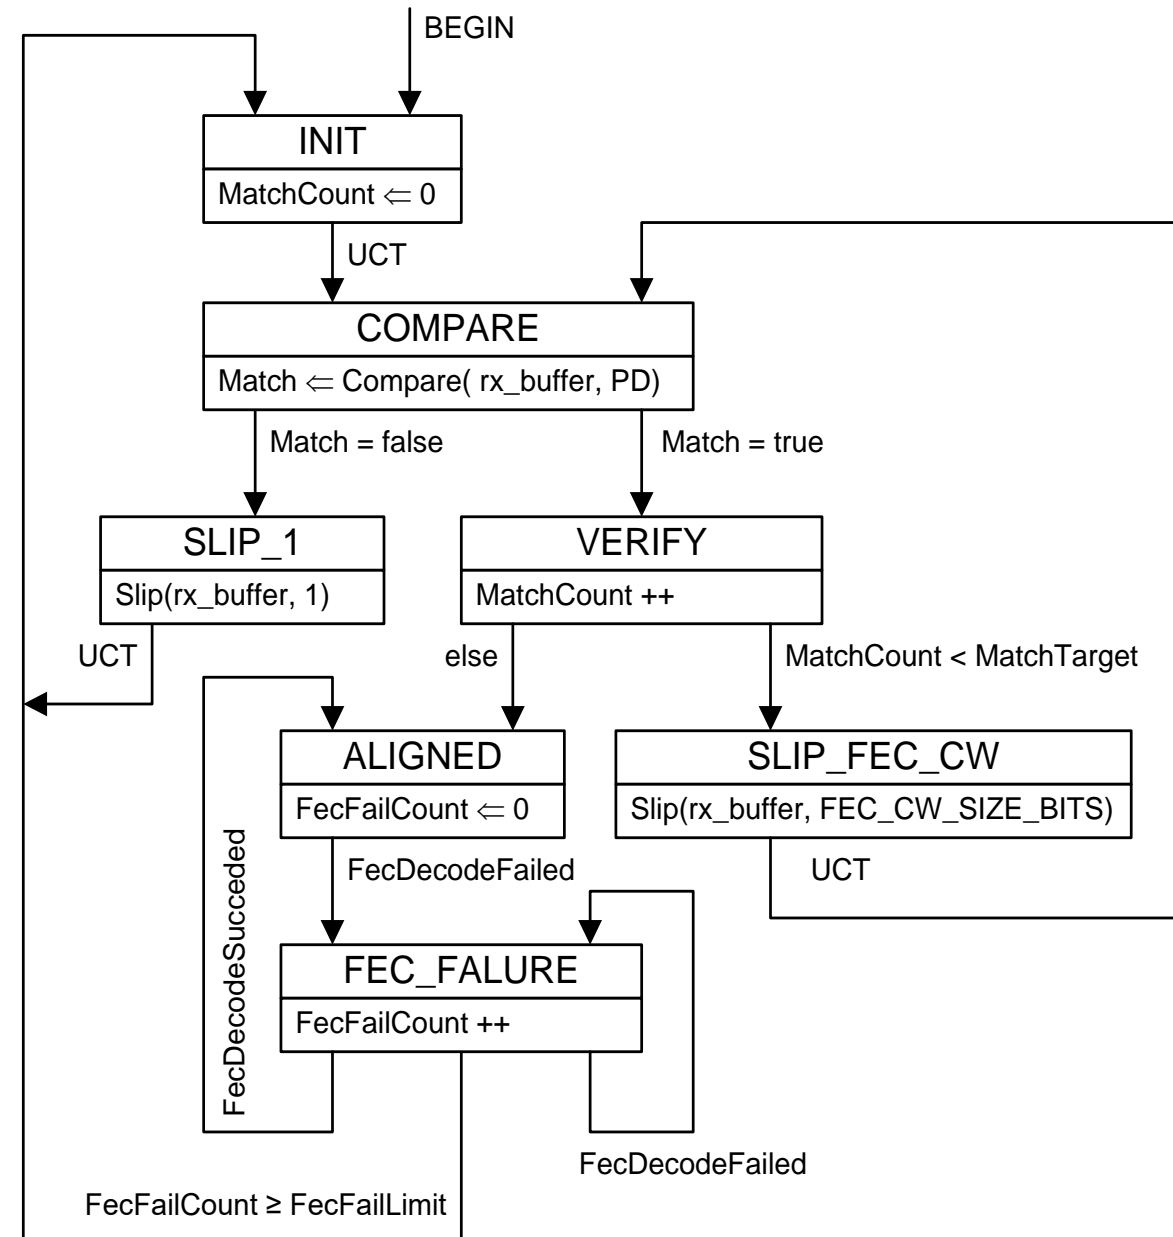
- For the receiver to align to the FEC codewords, it is enough to only mark the first parity block with a special **Parity Delimiter (PD)**.
- The PD should start with the special code-point 0-1111 plus 6 additional bits to make it as close to balanced code as possible (i.e., 5 zeroes, 6 ones)
 - 011-1100-1010_b (or **0x3ca**) ← Wow! A divine sign!!! Go ahead, argue with this proposal ()
- If we use the code-point to mark the first parity block, then in this block we only have $257-11=246$ bits available for the actual parity value.
- Subsequent 10 parity blocks are used to carry parity entirely, i.e., all 257 bits of each block.
- The total size of parity is $246 + 10 \times 257 = 2816$ as planned.

ONU FEC alignment mechanism

- Downstream never needs to shorten the codewords. Every codeword has a fixed size of 72 257b blocks (payload + parity)
- After reboot, the ONU needs to align to FEC codewords
 - Direct step from bit alignment to FEC CW alignment
 - There is no 257b block alignment in between. (257b blocks are not alignable; requires additional alignment markers)

To do the alignment, the ONU does the following:

1. Hunt for PD (11-bits) in the received bit stream
 - a) If PD is not matched, $MatchCount = 0$, slip incoming stream by 1 bit, go to (1).
 - b) If PD is matched, $MatchCount ++$, slip incoming stream by one full FEC codeword (18504 bits), go to (1)
2. When $MatchCount$ reaches a $MatchTarget$ threshold (5?), declare the alignment
3. If FEC decoder consistently indicates decoding failure for $FecFailLimit$ (3?) consecutive FEC CWs, set $MatchCount = 0$, go to 1 to realign.



Reliability of Alignment Mechanism

- The alignment SD will often pursue false leads before stumbling on the true delimiter. That is OK and expected.
- Because the alignment mechanism relies on the feedback from the FEC Decoder, probability of a wrong alignment is 0. The ONU will get aligned eventually. The question is how long that would take?
- The efficiency of FEC alignment depends on the choice of
 - *PD matching Hamming threshold*
 - *MatchTarget*
 - *FecFailLimit*
- We should find the best combination of these values to minimize the alignment time

Simulation study

- 1 million alignment attempts for each configuration
- Input BER = 0.01
- Match Target = 5 (i.e., need to match five delimiters in a row to declare alignment)
- *Hamming threshold* = {0, 1, 2}
- *MatchTarget* = 5 (i.e., need to match five delimiters in a row to declare alignment)
- *FecFailLimit* = 3
 - Time to detect loss of alignment (persistent FEC decode failure) = 10 μ s (this includes Rx time for 3 FEC CWs and FEC decoder delay)
 - This parameter was not relevant as alignment to a wrong position was not observed in the simulation




Simulation Results for varying match hamming threshold

False Leads




Missing True PD

Time to Alignment (μs)

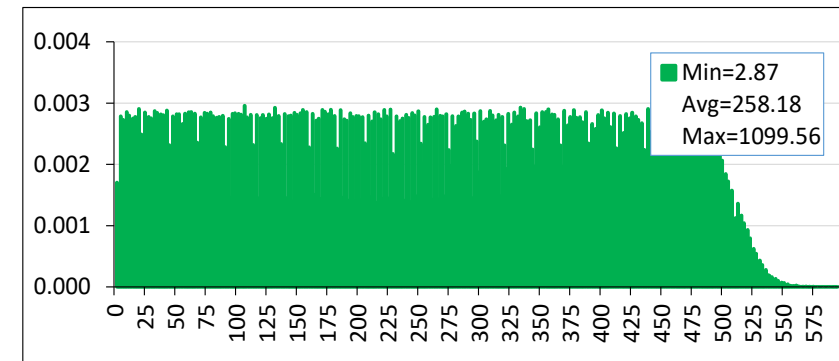


Hamming = 0



Hamming = 1




Hamming = 2



Conclusion

- If we hunt for a PD with $H=2$, the Alignment algorithm will pursue many more false leads. But the true PD position is never missed (in 1M tries, anyway).
- If we hunt for a PD with $H=0$, the algorithm will pursue very few false leads. But in about 42% of tests, the true delimiter will be missed the first time. Sometimes the true PD position may be missed multiple times (max observed in 1 M tests is 16 times)
- However the main metrics for the performance of the alignment algorithm is the total time it takes an ONU to align. Simulations show that using $H=0$ provides by far the shortest alignment time.
 - Average: 12 75 μ s
 - Maximum: 147.68 μ s
- Considering the typical interval between Discovery GATEs of 1 second, the alignment time (in any configuration) consumes a negligible fraction of the overall ONU Discovery and Registration time.



Backup



Hamming Distance Threshold

- When we hunt for the PD in the input stream, we should allow a match with some Hamming distance (H = 0? 1? 2?)
- Probability of a random block (not PD) looking like PD

$$\text{if } H = 2, \quad P(\text{falsePD}) = \frac{\binom{11}{2} + \binom{11}{1} + 1}{2^{11}} = 0.0327$$

$$\text{if } H = 1, \quad P(\text{falsePD}) = \frac{\binom{11}{1} + 1}{2^{11}} = 0.00586$$

$$\text{if } H = 0, \quad P(\text{falsePD}) = \frac{1}{2^{11}} = 0.000488$$

Expected Number of Matches

- Expected number of false PDs matchings in one FEC CW (continuous scan):
 - if $H = 2$, $F = 18504 \times 0.0327 = 605$
 - if $H = 1$, $F = 18504 \times 0.00586 = 108$
 - if $H = 0$, $F = 18504 \times 0.000488 = 9$
- On average, the comparator will encounter half the expected number of false PDs.
- For each false PD, the comparator will usually try one more position exactly 18504 bits away, before resuming the hunt.

Average Time to Align

- The average time before arriving to the true PD position is therefore $740ns \times F/2$
 - if $H = 2$, $Time = 223 \mu s$
 - if $H = 1$, $Time = 40 \mu s$
 - if $H = 0$, $Time = 3.3 \mu s$

Missing True PD

- Once the scanner arrived to the proper PD position, what is the probability of missing it?
- To declare alignment, the comparator needs to match the PDs with Hamming Distance H exactly 18504 bits apart M times ($M = 5$?)
- Probability of alignment failure = probability that at least one of the comparisons will not match, i.e., at least one out of M PDs will have more than H errors.

$$P(PDmatch) = \sum_{h=0}^H \binom{11}{h} BER^h \times (1 - BER)^{11-h}$$

- Probability of mismatch of 1 or more PD out of M :

$$P(PDmismatch) = 1 - P(PDmatch)^M$$

Probabilities of missing true PD

H	Prob. of matching a single true PD P(PDmatch)	Prob. of mismatching at least one PD out of M P(PDmismatch), M=		
		3	4	5
0	0.895338254	0.282269	0.357388	0.424645
1	0.994820283	0.015459	0.020558	0.025632
2	0.999844627	0.000466	0.000621	0.000777