



400GBASE-ZR PCS/PMA Baseline Proposal

Ilya Lyubomirsky, Jamal Riani, Ben Smith, Arash Farhoodfar, Sudeep Bhoja, Inphi Corp., Rich Baca, Microsoft Corp., Gary Nicholl, Cisco Systems, Inc., Mike Sluyski, Acacia Communications, Inc.

IEEE P802.3cn Meeting, Long Beach, CA, Jan. 14-15, 2019

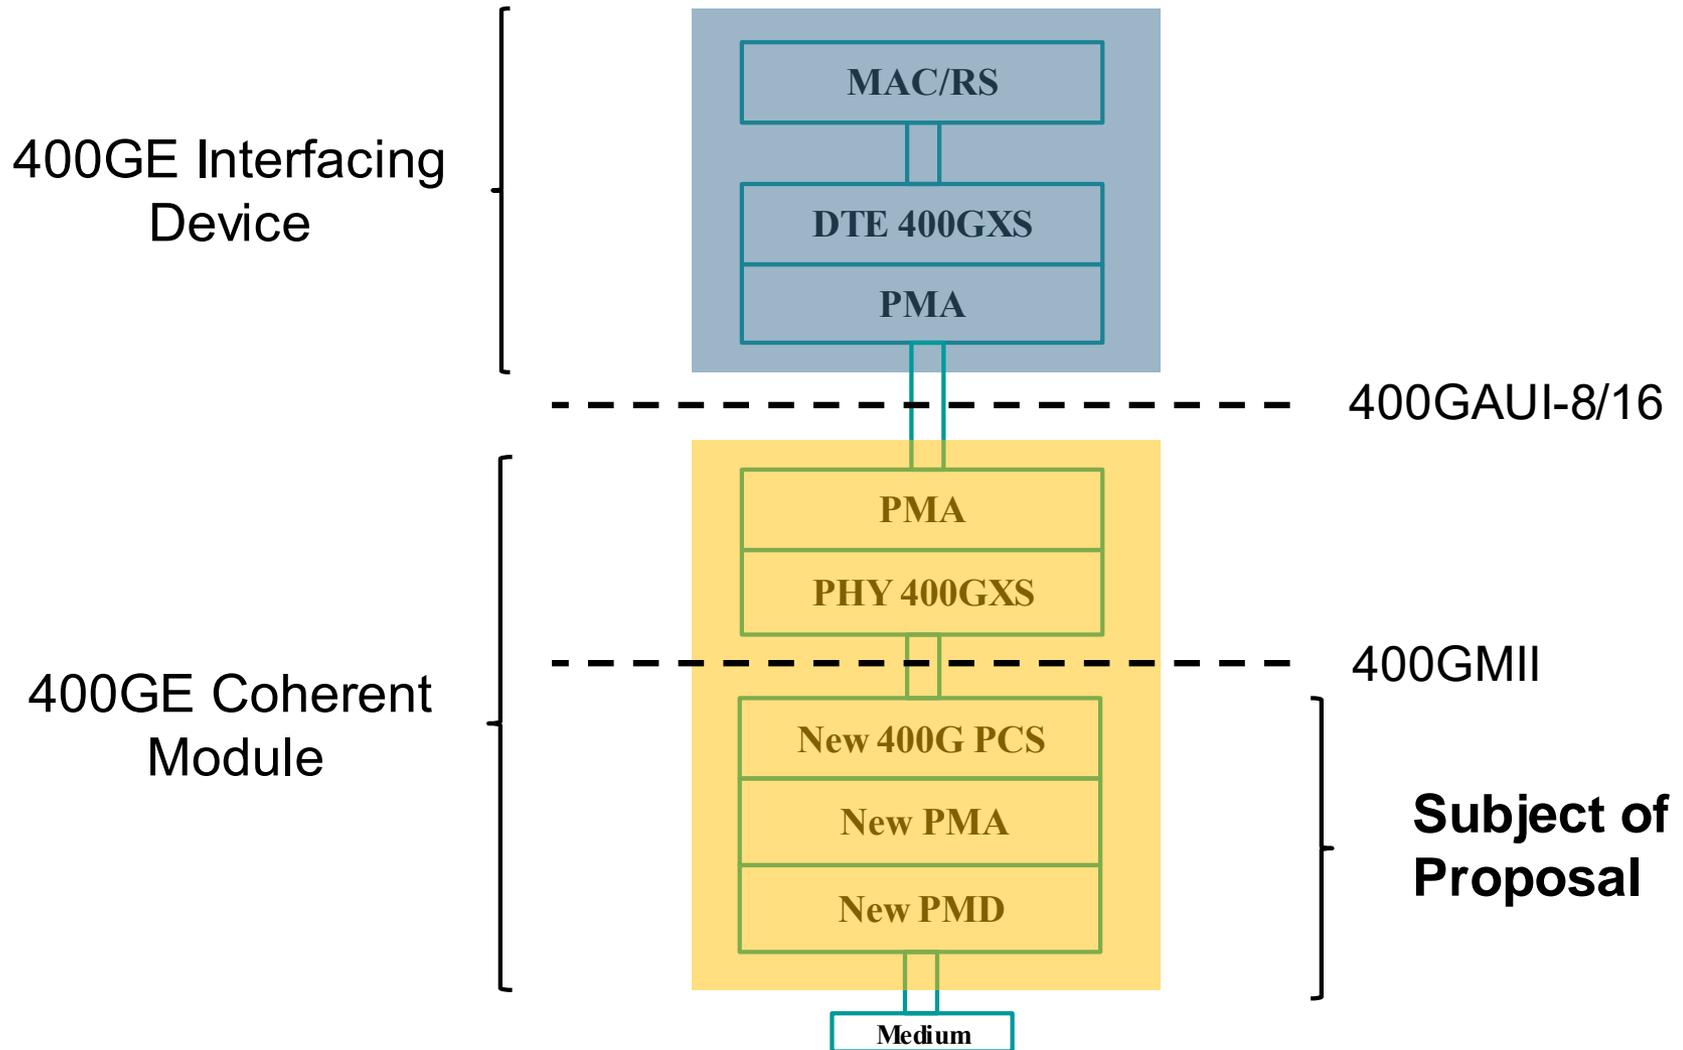
Supporters

- Pete Anslow, Ciena
- Paul Brooks, Viavi Solutions
- John DeAndrea, Finisar
- Mark Gustlin, Cisco
- Tad Hofmeister, Google
- Mark Nowell, Cisco
- Rob Stone, Broadcom
- Steve Trowbridge, Nokia
- Tomoo Takahara, Fujitsu
- Matt Brown, Macom
- Jean-Michel Caia, Fiberhome
- Mark Filer, Microsoft
- Steve Gorshe, Microchip
- Jeffery Maki, Juniper
- Matthew Schmitt, Cablelabs
- Peter Stassar, Huawei
- Chengbin Wu, ZTE
- Chongjin Xie, Alibaba

Summary

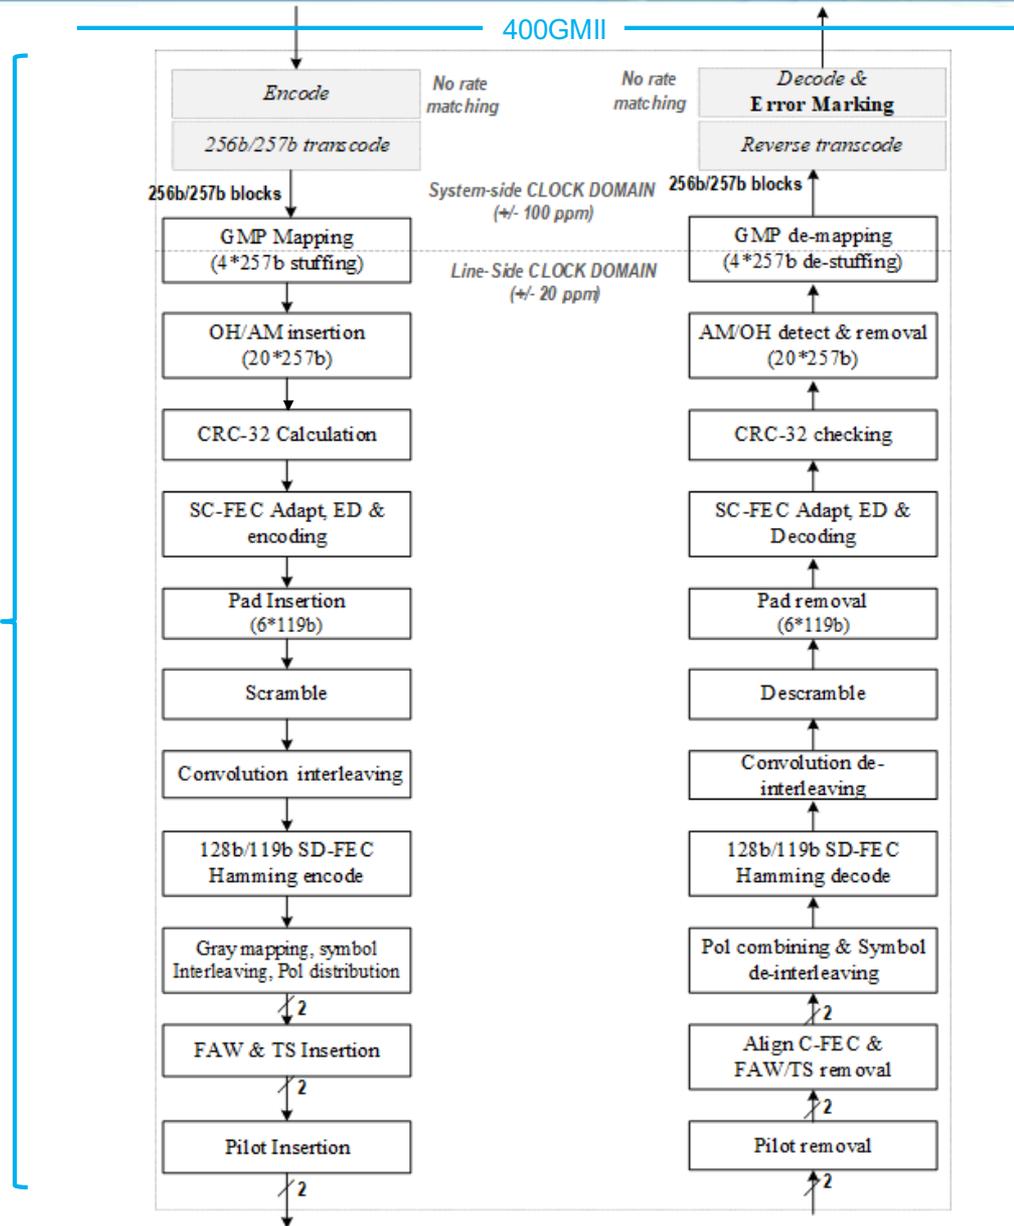
- OIF 400ZR 0.10 Draft specs. have been provided by OIF to IEEE 802.3cn via the liaison letter of Nov. 8, 2018
- At Nov. 2018 meeting, the 802.3cn Task Force adopted the OIF 400ZR modulation format (DP-16QAM) and FEC (CFEC) for 400GBASE-ZR; straw poll showed strong support (38 to 0) for adopting OIF 400ZR framing
- This presentation proposes the OIF 400ZR framing format for 400GBASE-ZR PCS and PMA, providing sufficient detail for an initial draft baseline

High Level View

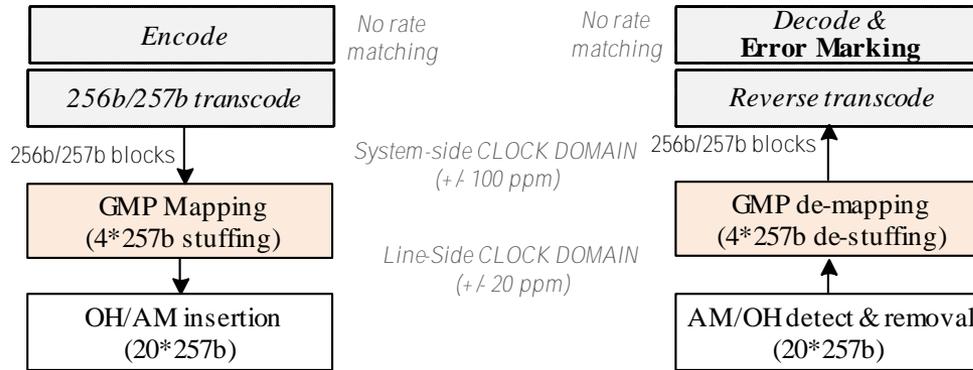


400ZR PCS and PMA Architecture Proposal

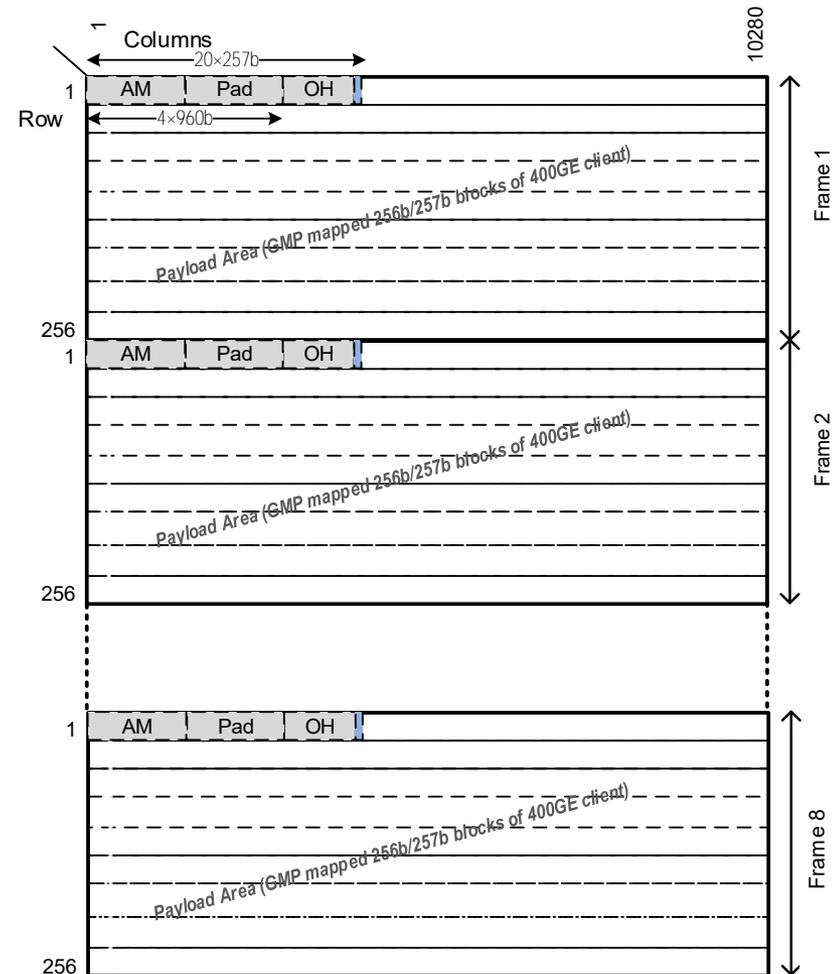
PCS/PMA



GMP Mapping to 400ZR Frame

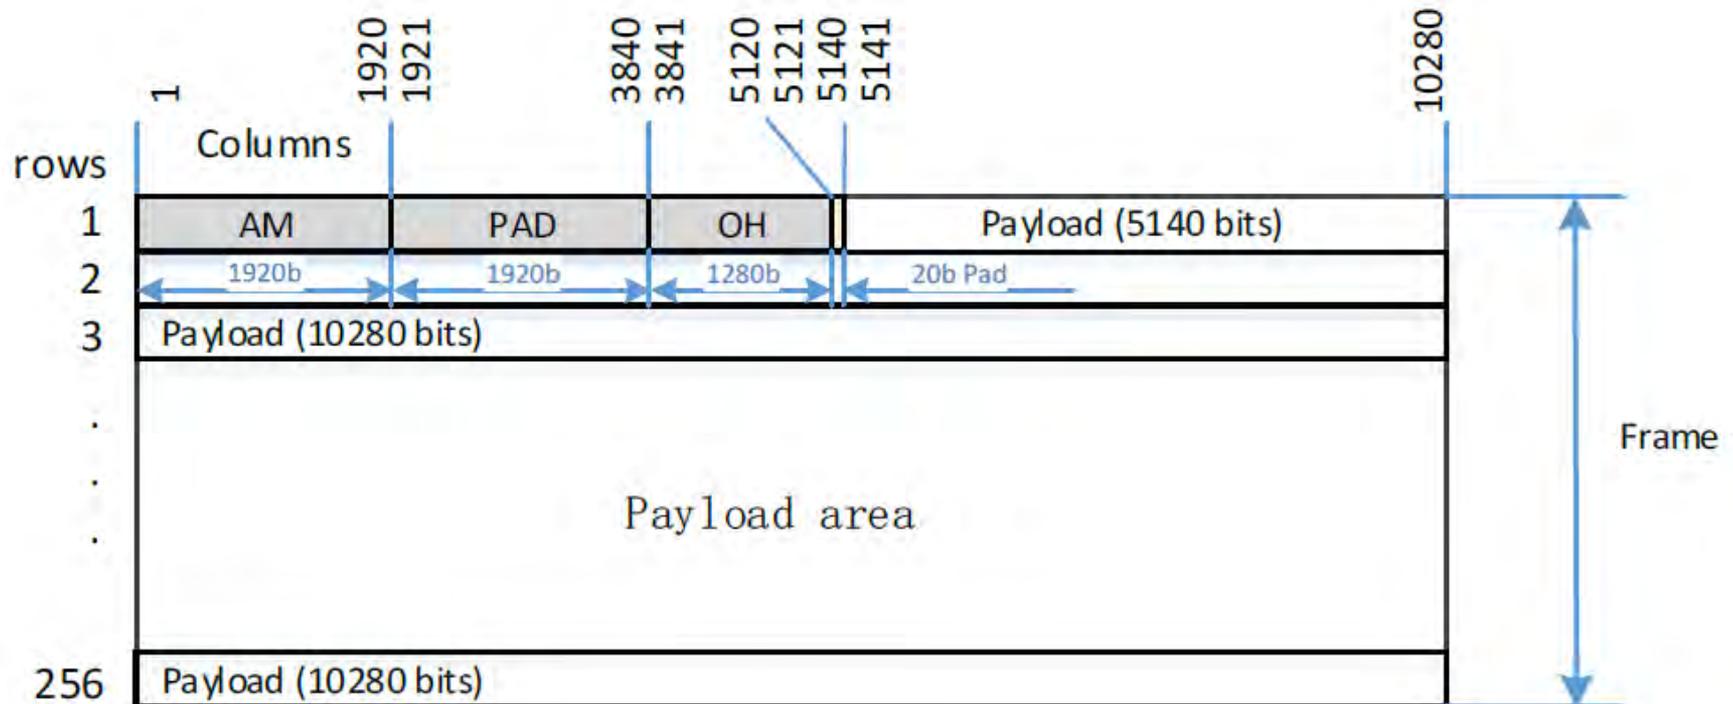


400ZR Frame Structure



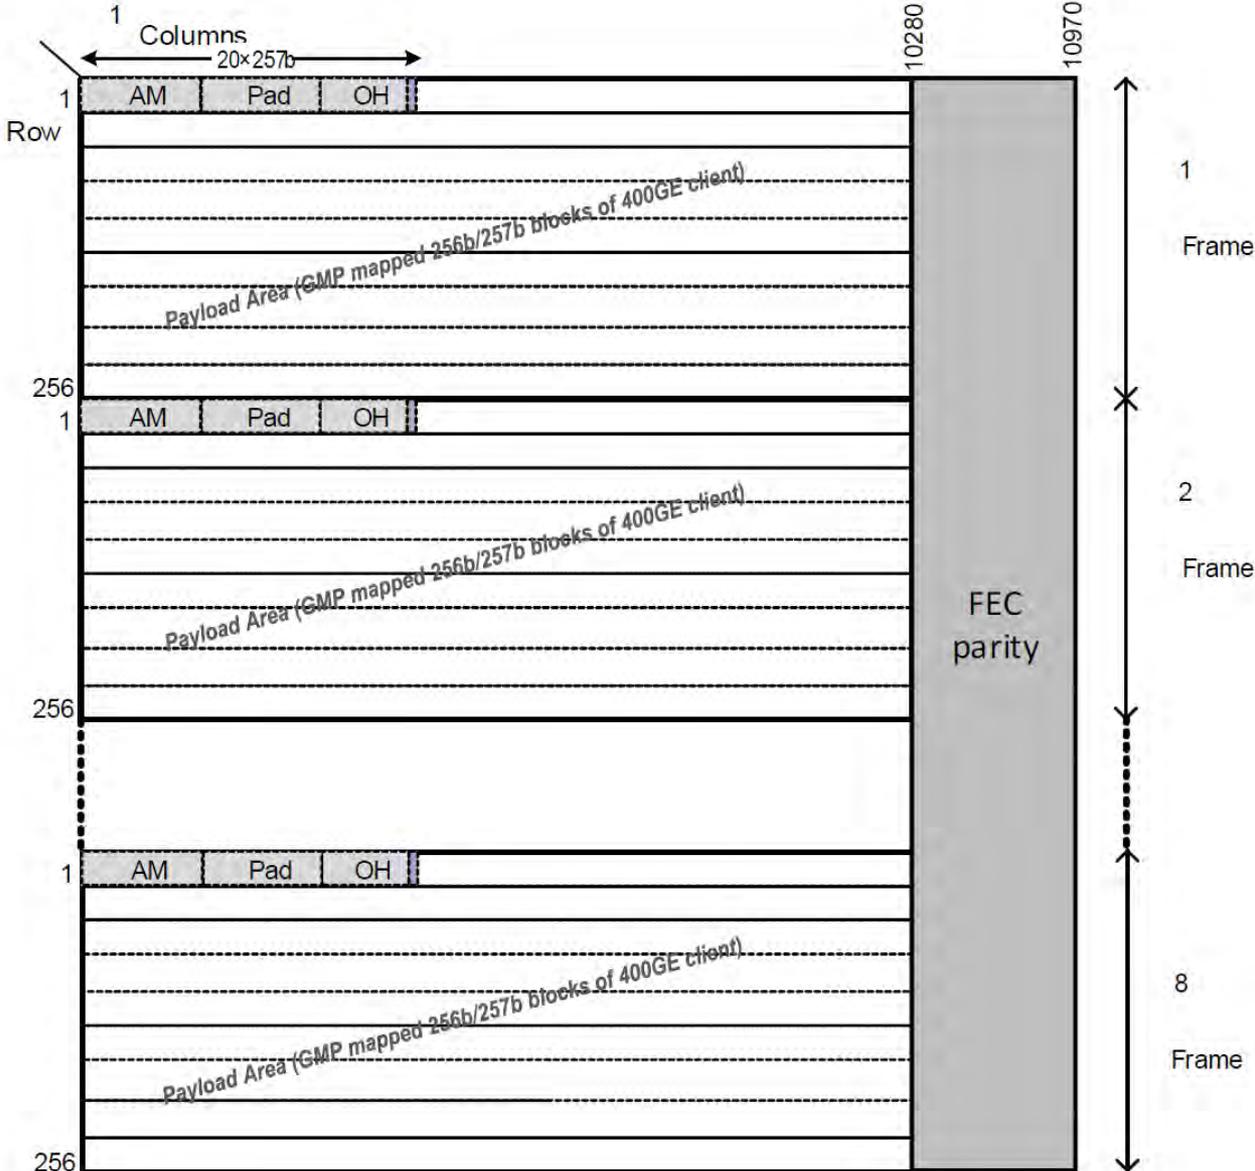
- 400GE signal is mapped to a 400ZR frame as a 256b/257b block stream
- GMP mapping (4x257b) is used to rate-adapt payload to local reference with +/- 20 ppm clock accuracy

ZR Frame Structure

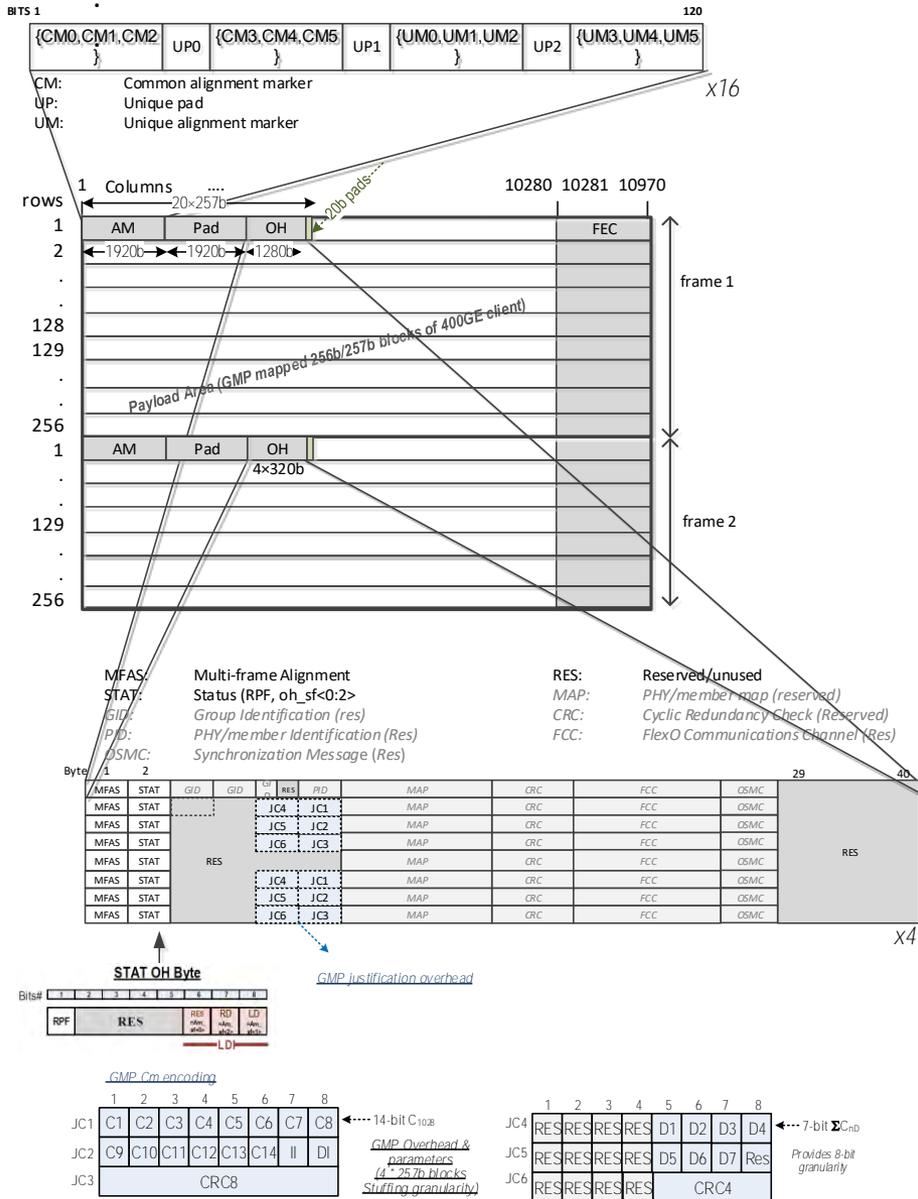


- 400ZR frame uses the 16×120-bit AMs defined by clause 119 [IEEE 802.3bs]
- Immediately following the 1920 bit AM are 1920 bits of all-Zero's PAD
- 1280 bit ZR OH area includes GMP mapping bytes
- An additional 20b PAD follows the ZR OH to provide 257b block alignment in the payload area

ZR Multi-Frame Structure



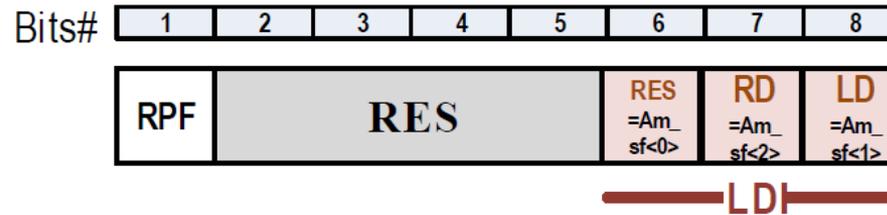
ZR Frame Overhead



- GMP Overhead (JC Bytes OH) is carried once per GMP payload envelop (once per 4-frame multiframe)
- It carries the encoded 14-bit Cm(t) (i.e. 4×257b block count value) in C1-14 bits of JC1&JC2 (C1 = MSB, ..., C14= LSB) and the encoded 8-bit SC_{ND}(t) (cumulative value of C_{ND}(t)) in D1-D7 bits of JC2&JC3 (D1=MSB,..., D7 = LSB)
- C_m(t) is protected by a CRC8 (carried in JC3 OH byte) and SC_{ND}(t) is protected by a CRC4 (carried in the four LSBs of JC6 OH byte)
- JC3 OH CRC8 calculation is described in ITU-T G.709 Annex D.3; it uses the generator polynomial $g(x) = x^8 + x^3 + x^2 + 1$; JC6 OH CRC4 calculation uses the generator polynomial $g(x) = x^4 + x + 1$
- Multi-frame Alignment Signal (MFAS) overhead byte is duplicated in all four 320-bit OH instances. It is incremented in every 400ZR frame from 0x00 to 0xFF, and provides a 256-frame multi-frame sequence following ITU-T G.709.1 Clause 9.2.1
- STAT overhead byte is present in every 400ZR frame, but only carried in the first of the four 320-bit OH instances. It includes the 1-bit RPF following the definition in ITU-T G.709.1 Clause 9.2.5.1, and 3-bit LDI fields following IEEE 802.3bs

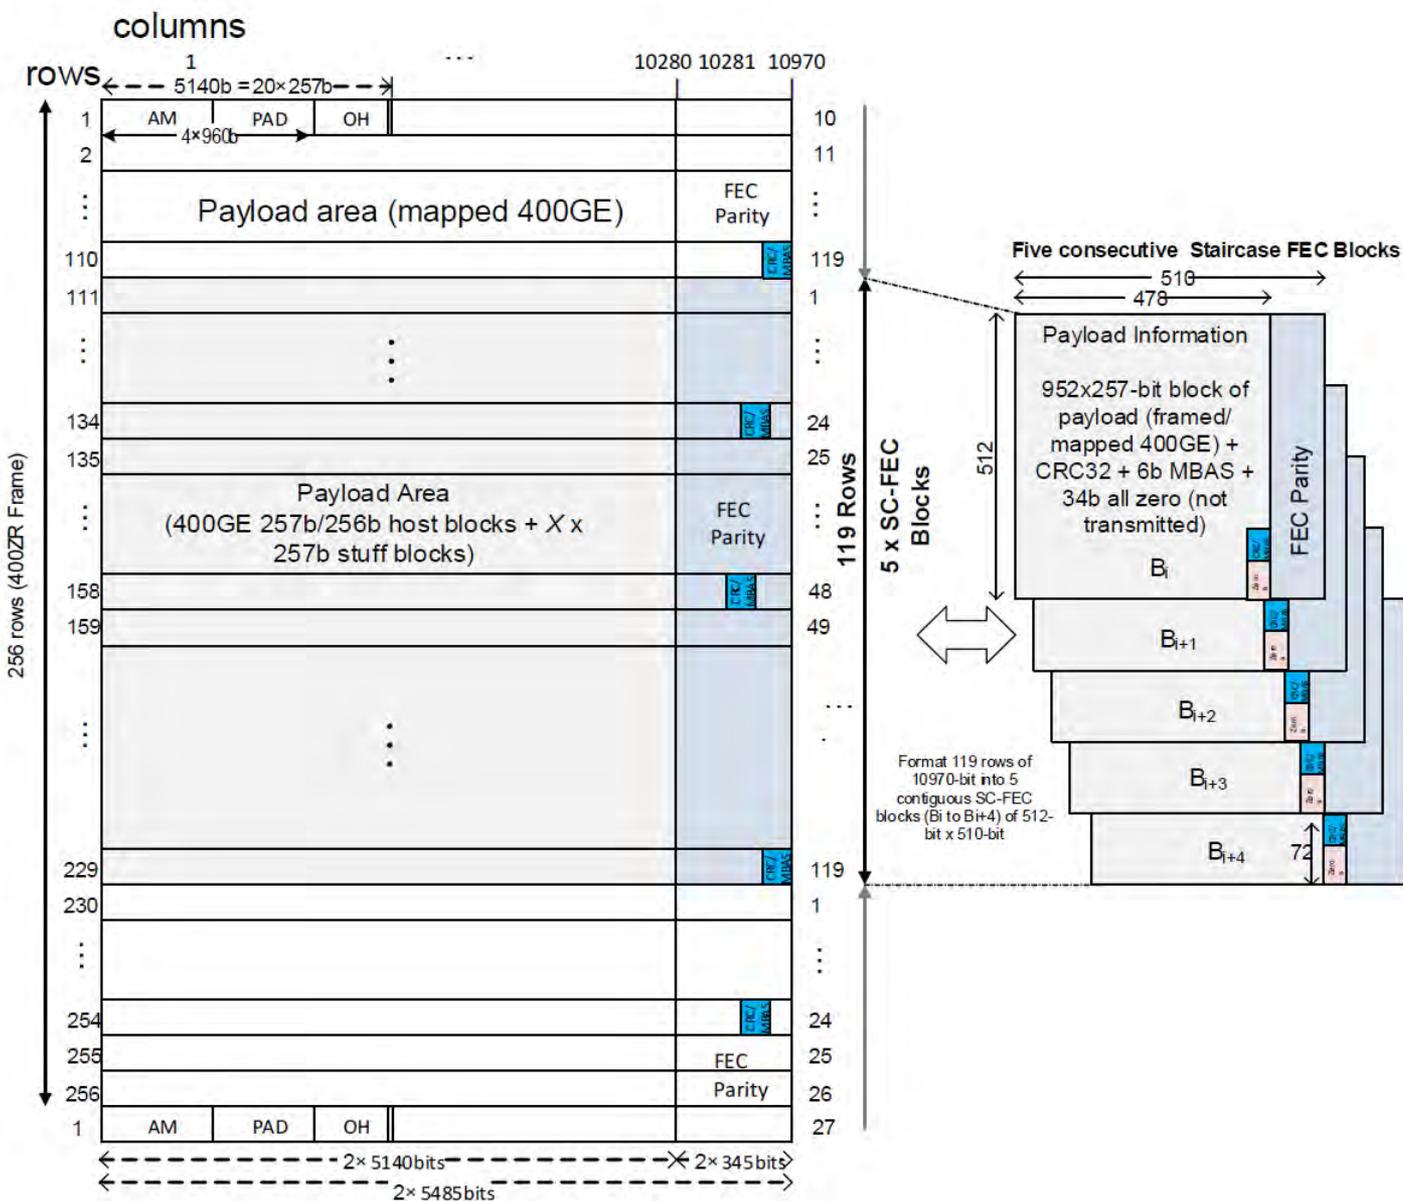
Link Status Monitoring and Signaling (STAT)

STAT OH Byte

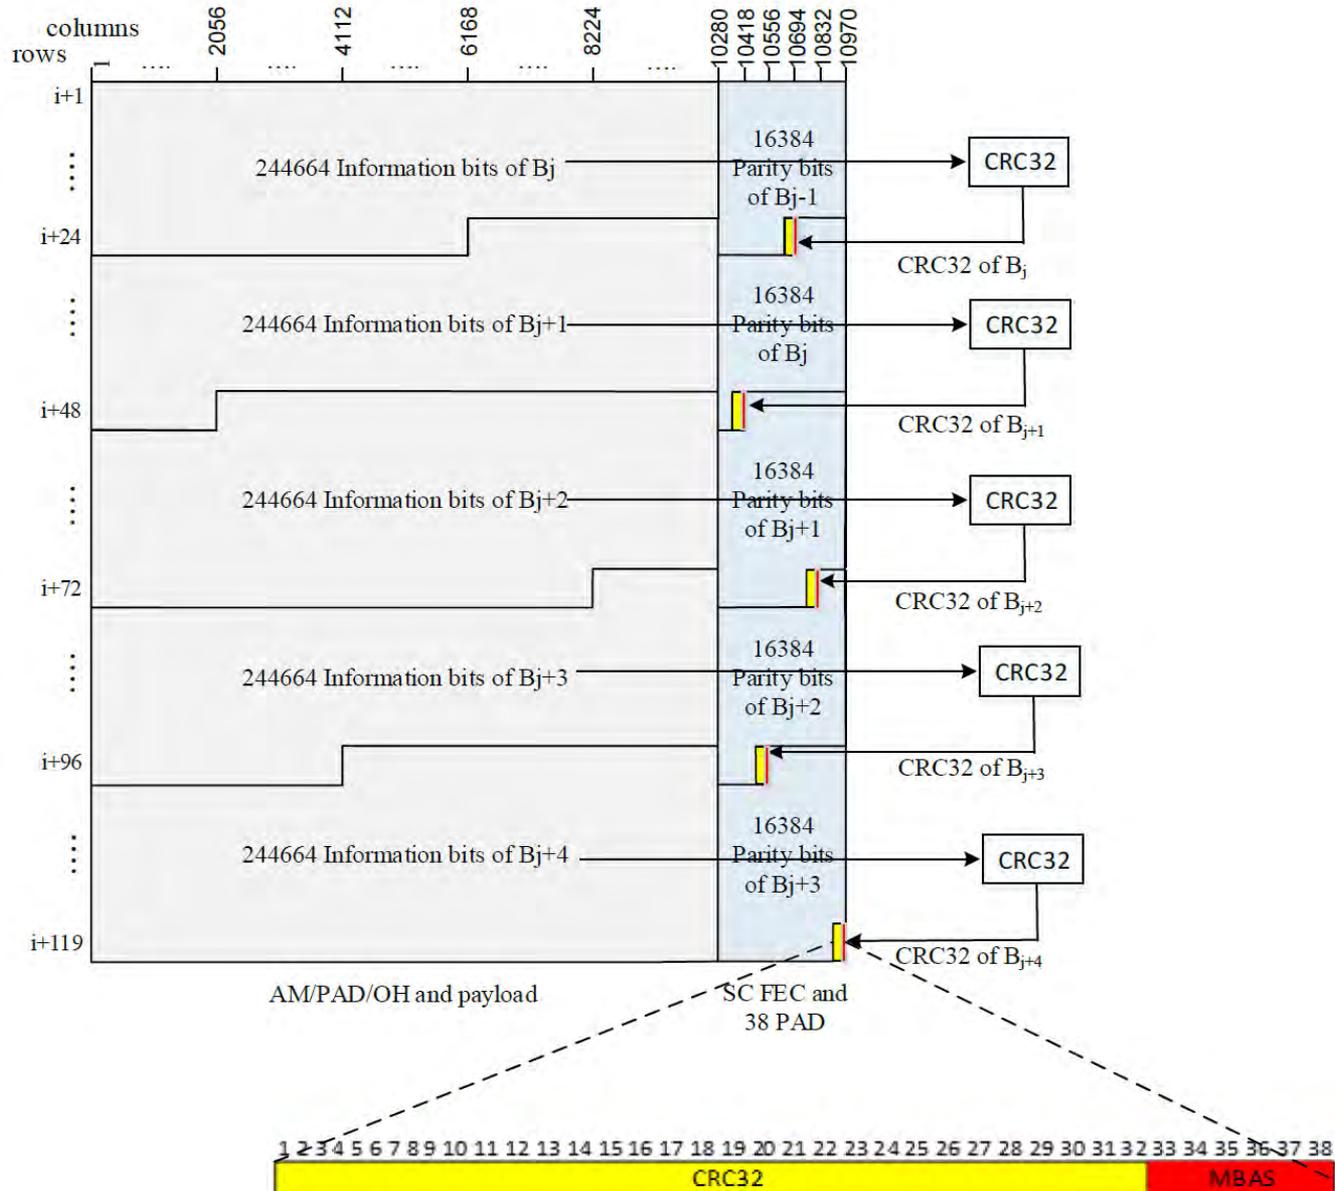


- The Remote PHY Fault (RPF) bit indicates signal fail status detected at the remote 400ZR sink function in the upstream direction, and follows ITU-T G.709.1 Clause 9.2.5.1 definition. RPF is set to "1" to indicate a remote 400ZR PHY defect indication; otherwise, it is set to "0". The RPF field is located in bit 1 of the STAT field
- Link Degrade Indication (LDI) follows the IEEE 802.3bs definition, which specified three bits in the AM field (am_sf<2:0>) to carry the LDI. Bit am_sf<2> is defined as a Remote Degrade (RD) signal, bit am_sf<1> is defined as a Local Degrade (LD) signal and bit am_sf<0> is reserved. The information in am_sf<0> shall be carried in STAT overhead bit 6. The status information in am_sf<2> shall be carried in STAT overhead bit 7. The status information in the LD (am_sf<1>) bit shall be carried in STAT overhead bit 8

ZR Frame to SC-FEC Mapping



CRC-32 and MBAS



CRC-32

- A 32-bit CRC is calculated over the 244,664 input bits using the generator polynomial:
$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$
- Mathematically, the CRC value corresponding to the 244,664 input bits is defined by the following algorithm:
 - a) The first 32 bits of the frame are complemented.
 - b) The 244,664 bits of the protected fields are considered to be the coefficients of a polynomial $M(x)$ of degree 244,663. (The first bit of the 244,664 input bits corresponds to the $x^{244,663}$ term and the last bit of the 244,664 input bits corresponds to the x^0 term)
 - c) $M(x)$ is multiplied by x^{32} and divided (modulo 2) by $G(x)$, producing a remainder $R(x)$ of degree ≤ 31 .
 - d) The coefficients of $R(x)$ are considered to be a 32-bit sequence.
 - e) The bit sequence is complemented and the result is the CRC.
- The 32 bits of the CRC value are placed with x^{31} term as the left-most bit of the CRC32 field and the x^0 term as the right most bit of the CRC32 field. (The bits of the CRC are thus transmitted in the order: $x^{31}, x^{30}, \dots, x^1, x^0$)
- The 6 bit MBAS (see next slide) is appended after the 32-bit CRC

Multi-Block Alignment Signal (MBAS)

- To synchronize the state of the Error Decorrelator (ED) controllers between the receiver and the transmitter, the SC-FEC scheme uses a 7-bit SC FEC Multi Block Alignment Signal (MBAS) which provides a 128 block sequence.
- The six most significant bits of the 7-bit MBAS are transferred between source and sink in the 6-bit MBAS overhead, which is located in bits 33 to 38.
- The numerical value represented in the six MBAS overhead bits will be incremented every two SC-FEC blocks

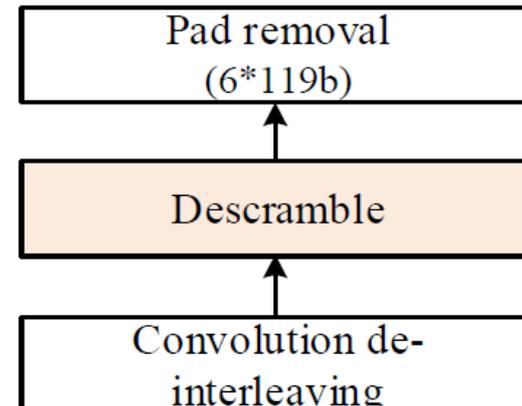
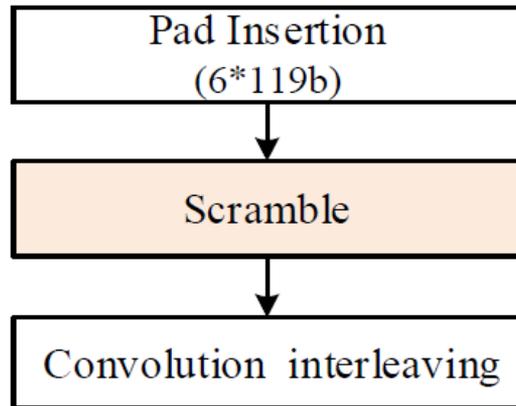
7-bit Multi-Block Alignment Signal							
	1	2	3	4	5	6	7
6-bit MBAS OH							
	33	34	35	36	37	38	
	:						
0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	1
2	0	0	0	0	0	1	0
3	0	0	0	0	0	1	1
4	0	0	0	0	1	0	0
5	0	0	0	0	1	0	1
6	0	0	0	0	1	1	0
7	0	0	0	0	1	1	1
	:						
	:						
125	1	1	1	1	1	0	1
126	1	1	1	1	1	1	0
127	1	1	1	1	1	1	1
0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	1
2	0	0	0	0	0	1	0
	:						
	:						

SC-FEC



- The SC-FEC is a 512-bit x 510-bit generalized staircase code that works in conjunction with an error de-correlator
- SC-FEC encoding/decoding processes shown above are defined by reference to ITU-T G.709.2/Y.1331.2 (2018), OTU4 long-reach interface, Annex A

Frame Synchronous Scrambler

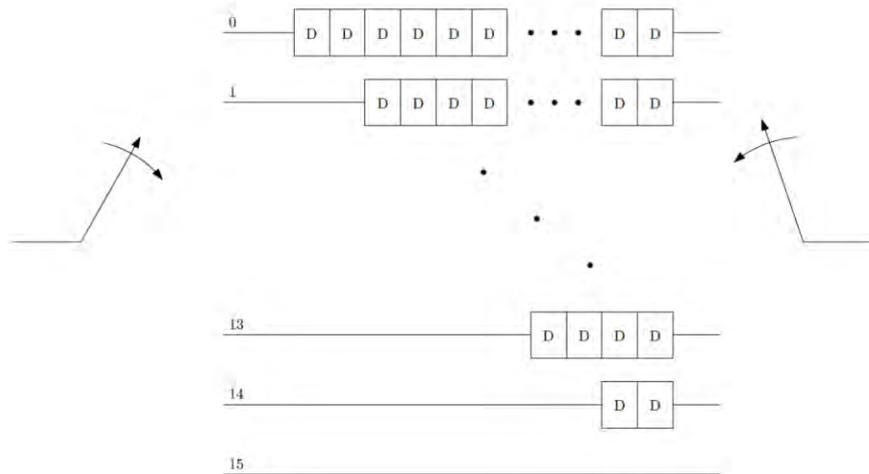


- The operation of the scrambler shall be functionally equivalent to that of a frame-synchronous scrambler with generating polynomial:

$$x^{16} + x^{12} + x^3 + x + 1$$

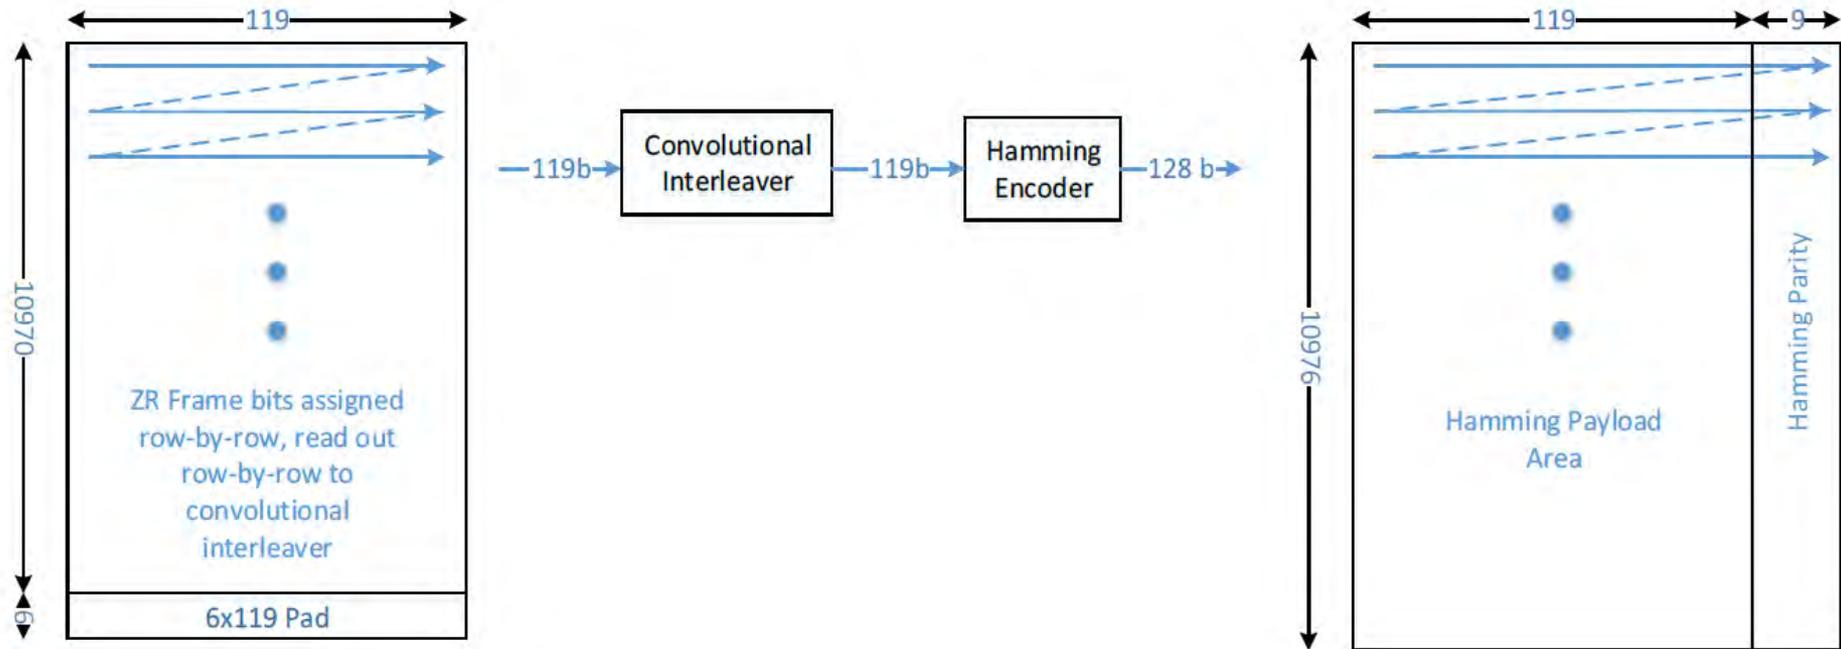
- The scrambler/descrambler resets to 0xFFFF on row 1, column 1 of the five SC-FEC block structure and subsequent 714-bit (6 x 119b) pad insertion and the scrambler state advances during each bit of the 5xSC-FEC blocks. In the source function, all payload bits (included SC-FEC parity) are scrambled. At the sink the scrambler is synchronized (initialized) at the start of each payload

Convolutional Interleaver



- Each delay operator “D” represents a storage element of 119b. From one row to the next lower row, two delay operators are deleted
- At time i , the input and output switches are aligned at row b_i :
 - 1) A block of 119b is read from row b_i
 - 2) The contents of row b_i are shifted to the right by 119b
 - 3) A block of 119b is written to row b_i
 - 4) The switch position is updated to $b_{i+1} = b_i + 1 \pmod{16}$
- Initialization of the convolutional interleaver is defined to occur at the start of every DSP Super frame, which contains 5 SC-FEC blocks

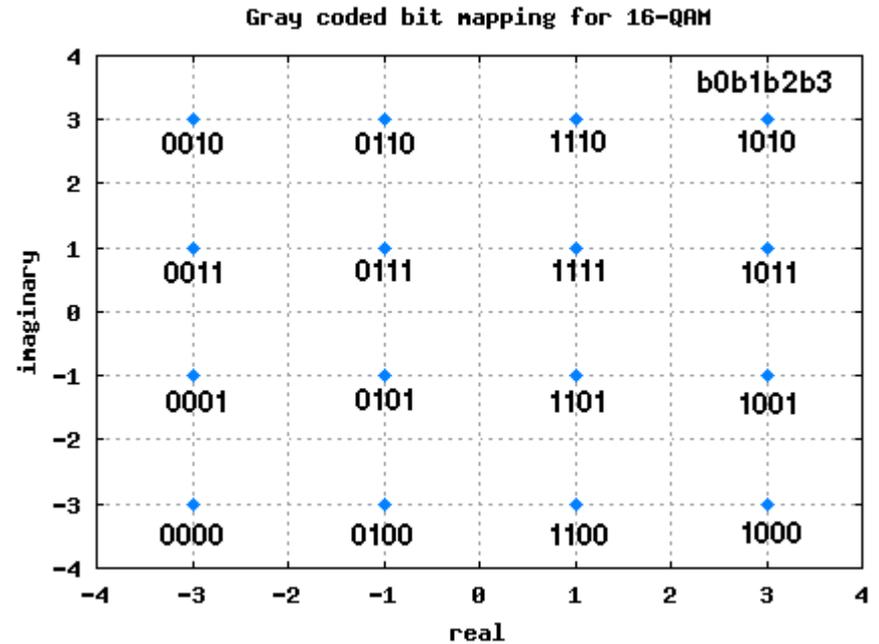
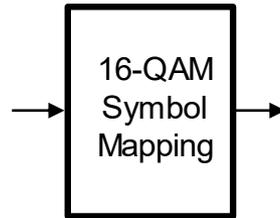
Hamming Code



- The 119b outputs of the convolutional interleave are encoded by a systematic (128,119) double-extended Hamming code
- With soft decision decoding, the (128,119) Hamming code increases the NCG from 9.4 dB to ~10.8 dB with ~7.56% added overhead

16-QAM Symbol Mapping

100110101000...



Bits are Grey mapped to 16-QAM symbols

Forming DP-16QAM Symbols

- Each 128-bit code word is mapped to 16 DP-16QAM symbols,

$$S = [s_0, s_1, \dots, s_{15}],$$

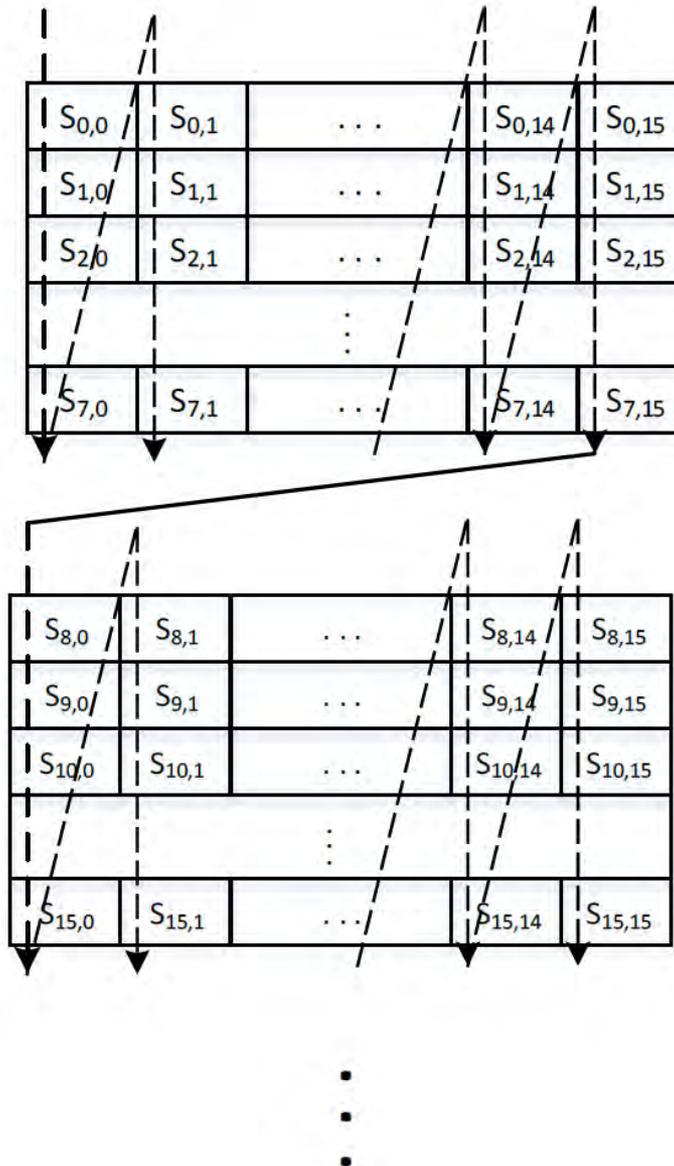
where for $i = 0$ to 15

- (c_{8i}, c_{8i+1}) maps to the in-phase (I) component of the X-pol of s_i
 - (c_{8i+2}, c_{8i+3}) maps to the quadrature-phase (Q) component of the X-pol of s_i
 - (c_{8i+4}, c_{8i+5}) maps to the I component of the Y-pol of s_i
 - (c_{8i+6}, c_{8i+7}) maps to the Q component of the Y-pol of s_i
- In each signaling dimension, we define the following map from binary label to symbol amplitude:

$$(0,0) \rightarrow -3, (0,1) \rightarrow -1, (1,1) \rightarrow +1, (1,0) \rightarrow +3$$

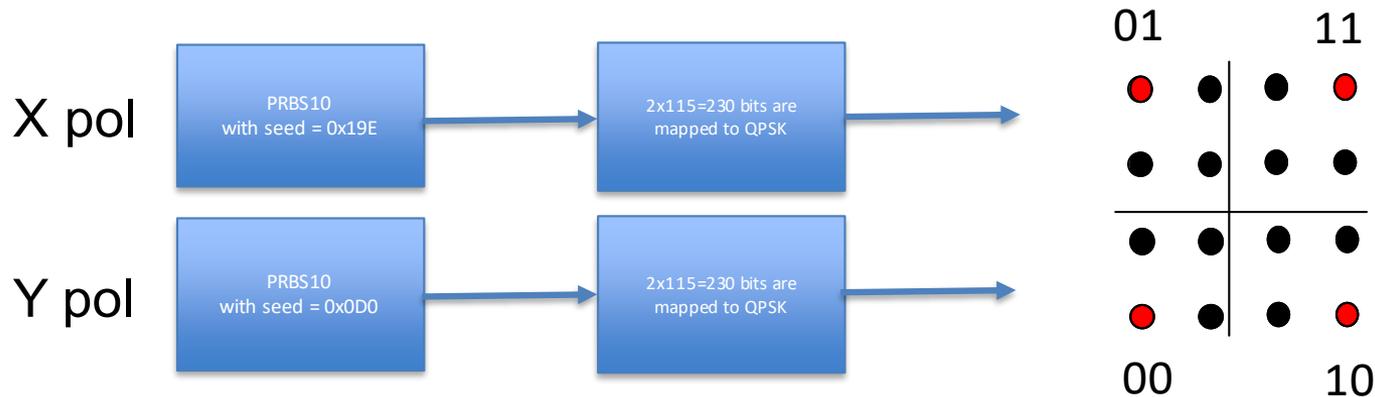
DP-16QAM Symbol Interleaving

8-way
interleaved
Hamming
Codewords



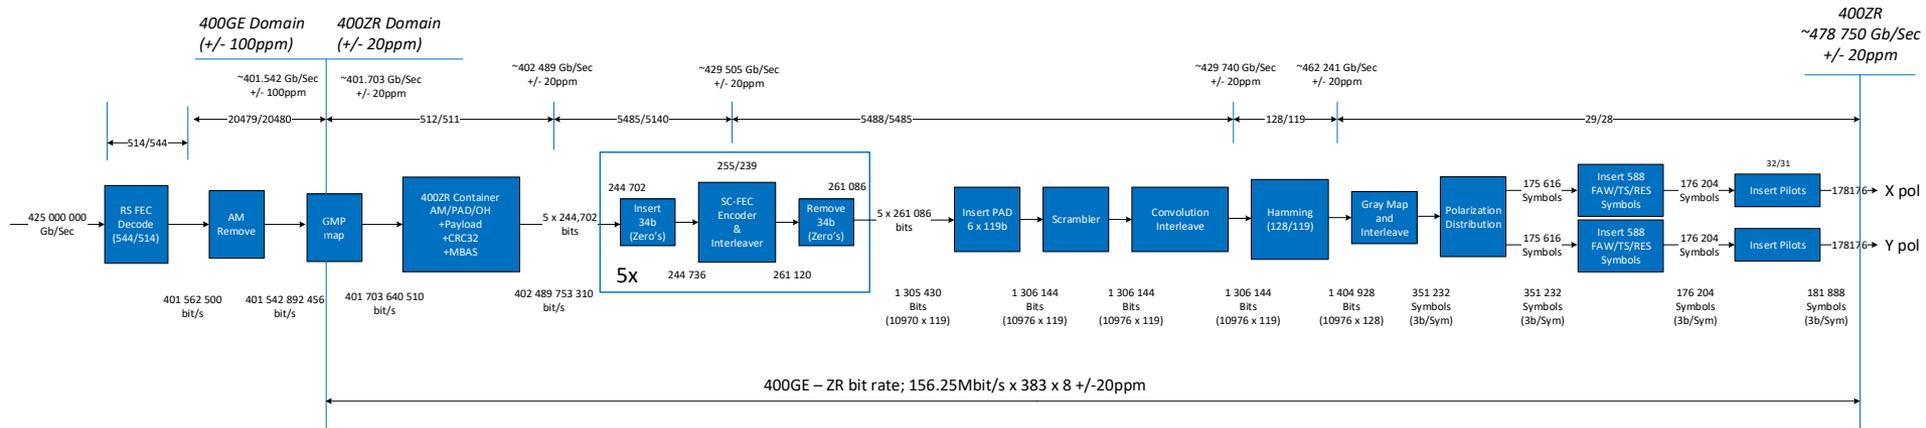
Pilot Symbols

Pilot symbols are added periodically to aid Rx DSP carrier phase recovery and enable absolute phase detection for better performance



- Pilot symbol inserted with a period of 32 QAM symbols
- Different pilot sequences used for X and Y polarizations

Frame Expansion Rate



GMP

FlexO/400ZR frame

	400GBASE-R Client-rate [bps]	After Client FEC Termination [bps]	After 400GE AM Removal [bps]	Before 400ZR AM/PAD/OH insert [bps]	Before SC-FEC + MBAS + CRC32 [bps]	Before [6 x 119b] pad insert [bps]	Before Hamming [bps]	Before FAW/TS/RES [bps]	Before Pilot Symbol insertion [bps]	400ZR Bit rate [bps]	400ZR Baud rate [bps]	
+100ppm	425 042 500 000	401 602 656 250	401 583 046 745	401 711 674 583	402 497 803 105	429 513 706 231	429 748 627 128	462 250 624 138	463 798 338 281	478 759 575 000	059 844 946 875	+20ppm
Nominal	425 000 000 000	401 562 500 000	401 542 892 456	401 703 640 510	402 489 753 310	429 505 116 129	429 740 032 328	462 241 379 310	463 789 062 500	478 750 000 000	059 843 750 000	Nominal
-100ppm	424 957 500 000	401 522 343 750	401 502 738 167	401 695 606 437	402 481 703 515	429 496 526 027	429 731 437 527	462 232 134 483	463 779 786 719	478 740 425 000	059 842 553 125	-20ppm

Recommendation

Adopt this proposal as baseline for the 400GBASE-ZR PCS and PMA