




802.3cz baseline proposal

Rubén Pérez-Aranda

Introduction and objectives

- Base on the link model presented in [1], link budget analysis were presented for 25 Gb/s and for 10 Gb/s (see [2], [3] and [4])
- These link budget analysis were based on a specific transmission scheme and a subset of parameters that define such scheme: PAM2, RS(544, 522) GF(2¹⁰), $F_s = (1.0625 \times \text{bit-rate})$ and DFE
- 25 Gb/s operation with low VCSEL bias current was demonstrated feasible for almost all evaluated VCSELs, in temperatures of 125°C and -40°C, over 40 m length OM3 cables, even with reduced EMB due to longer wavelengths, and 4 inline connections
- This contribution is a detailed **PCS** and **PMA** baseline proposal for **2.5, 5, 10, and 25 Gb/s** consistent with the presented link budget analysis
- This contribution does **NOT** include the needed extensions for optional **EEE**
- This contribution does **NOT** include **PMD**, **MDI** and **medium** specifications


Reference models



MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE
 25GMII = 25 GIGABIT MEDIA INDEPENDENT INTERFACE

Relationship of OMEGA PHYs to the ISO/IEC OSI reference model and the IEEE 802.3 Ethernet Model

PCS transmit ordering



PCS 64B/65B encoding (same as C/55)



Input Data	data ctrl header	Block Payload											
Bit Position:		0	1										64
Data Block Format:													
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ D ₇	0	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇				
Control Block Formats:		Block											
C ₀ C ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0x1E	C ₀	C ₁	C ₂	C ₃	C ₄	C ₅	C ₆	C ₇			
C ₀ C ₁ C ₂ C ₃ /O ₄ D ₅ D ₆ D ₇	1	0x2D	C ₀	C ₁	C ₂	C ₃	O ₄	D ₅	D ₆	D ₇			
C ₀ C ₁ C ₂ C ₃ /S ₄ D ₅ D ₆ D ₇	1	0x33	C ₀	C ₁	C ₂	C ₃		D ₅	D ₆	D ₇			
O ₀ D ₁ D ₂ D ₃ /S ₄ D ₅ D ₆ D ₇	1	0x66	D ₁	D ₂	D ₃	O ₀		D ₅	D ₆	D ₇			
O ₀ D ₁ D ₂ D ₃ /O ₄ D ₅ D ₆ D ₇	1	0x55	D ₁	D ₂	D ₃	O ₀	O ₄	D ₅	D ₆	D ₇			
S ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ D ₇	1	0x78	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇				
O ₀ D ₁ D ₂ D ₃ /C ₄ C ₅ C ₆ C ₇	1	0x4B	D ₁	D ₂	D ₃	O ₀	C ₄	C ₅	C ₆	C ₇			
T ₀ C ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0x87					C ₁	C ₂	C ₃	C ₄	C ₅ C ₆ C ₇		
D ₀ T ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0x99	D ₀					C ₂	C ₃	C ₄	C ₅ C ₆ C ₇		
D ₀ D ₁ T ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	1	0xAA	D ₀	D ₁					C ₃	C ₄	C ₅ C ₆ C ₇		
D ₀ D ₁ D ₂ T ₃ /C ₄ C ₅ C ₆ C ₇	1	0xB4	D ₀	D ₁	D ₂				C ₄	C ₅	C ₆ C ₇		
D ₀ D ₁ D ₂ D ₃ /T ₄ C ₅ C ₆ C ₇	1	0xCC	D ₀	D ₁	D ₂	D ₃				C ₅	C ₆ C ₇		
D ₀ D ₁ D ₂ D ₃ /D ₄ T ₅ C ₆ C ₇	1	0xD2	D ₀	D ₁	D ₂	D ₃	D ₄				C ₆ C ₇		
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ T ₆ C ₇	1	0xE1	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅			C ₇		
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ T ₇	1	0xFF	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆				


Figure 55-9—64B/65B block formats

Input Data	Sync	Block Payload										
Bit Position:		0	1	2								65
Data Block Format:												
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ D ₇	01	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇			
Control Block Formats:		Block Type Field										
C ₀ C ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	10	0x1e	C ₀	C ₁	C ₂	C ₃	C ₄	C ₅	C ₆	C ₇		
C ₀ C ₁ C ₂ C ₃ /O ₄ D ₅ D ₆ D ₇	10	0x2d	C ₀	C ₁	C ₂	C ₃	O ₄	D ₅	D ₆	D ₇		
C ₀ C ₁ C ₂ C ₃ /S ₄ D ₅ D ₆ D ₇	10	0x33	C ₀	C ₁	C ₂	C ₃		D ₅	D ₆	D ₇		
O ₀ D ₁ D ₂ D ₃ /S ₄ D ₅ D ₆ D ₇	10	0x66	D ₁	D ₂	D ₃	O ₀		D ₅	D ₆	D ₇		
O ₀ D ₁ D ₂ D ₃ /O ₄ D ₅ D ₆ D ₇	10	0x55	D ₁	D ₂	D ₃	O ₀	O ₄	D ₅	D ₆	D ₇		
S ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ D ₇	10	0x78	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇			
O ₀ D ₁ D ₂ D ₃ /C ₄ C ₅ C ₆ C ₇	10	0x4b	D ₁	D ₂	D ₃	O ₀	C ₄	C ₅	C ₆	C ₇		
T ₀ C ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	10	0x87					C ₁	C ₂	C ₃	C ₄	C ₅ C ₆ C ₇	
D ₀ T ₁ C ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	10	0x99	D ₀					C ₂	C ₃	C ₄	C ₅ C ₆ C ₇	
D ₀ D ₁ T ₂ C ₃ /C ₄ C ₅ C ₆ C ₇	10	0xaa	D ₀	D ₁					C ₃	C ₄	C ₅ C ₆ C ₇	
D ₀ D ₁ D ₂ T ₃ /C ₄ C ₅ C ₆ C ₇	10	0xb4	D ₀	D ₁	D ₂				C ₄	C ₅	C ₆ C ₇	
D ₀ D ₁ D ₂ D ₃ /T ₄ C ₅ C ₆ C ₇	10	0xcc	D ₀	D ₁	D ₂	D ₃				C ₅	C ₆ C ₇	
D ₀ D ₁ D ₂ D ₃ /D ₄ T ₅ C ₆ C ₇	10	0xd2	D ₀	D ₁	D ₂	D ₃	D ₄				C ₆ C ₇	
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ T ₆ C ₇	10	0xe1	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅			C ₇	
D ₀ D ₁ D ₂ D ₃ /D ₄ D ₅ D ₆ T ₇	10	0xff	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆			

Figure 49-7—64B/66B block formats

Trivial trans-coding from BASE-R used in serial interfaces while we save 1 bit

PCS transmit ordering – header




Binary additive scrambler (based on C/115)



- Same additive binary scrambler as C/115, same polynomial, different reset value.
- The 195840 bits composing the transmit block from the 36 RS codewords aggregation shall be scrambled prior to transmission using a binary scrambler that produces the same result as the implementation shown in figure. The shift register shall be initialized with a hexadecimal value of 0x02E57CA for each transmit block, where the leftmost digit corresponds to the initial value of register element $r[0]$. The initial value of $r[0]$ is xor-ed with the first bit from the first RS codeword to generate the first bit of the binary sequence that feeds the PAM2 mapper.
- Shift register is formally defined by the following MATLAB® code, where len is the length of the sequence to be generated ($len = 195840$), the variable out is the binary output sequence, and the variable $seed$ is the initialization value of the shift register ($seed = '02E57CA'$).
 - In the MATLAB® code, the shift register r is indexed from 1 to 25 due to language syntax constraints. $r(i)$, for i from 1 through 25 in the code corresponds to $r[i-1]$ of the figure.

```
function out = lfsr(len, seed)
% Reset
r = double(dec2bin(hex2dec(seed))) - double('0');
r = [zeros(1, 25-length(r)) r];
% Output generation
for i = 1:len
    out(i) = r(1);
    r = [mod(r(22)+r(25),2) r(1:24)];
end
end
```



PAM2 mapping, PMD mapping

- The PAM2 mapper shall map bits to symbols as:
 - bit = 0 \longleftrightarrow pam_symbol = -1
 - bit = 1 \longleftrightarrow pam_symbol = +1
- The PMD transmit function shall convert the symbol stream from the PAM2 mapper into optical symbols stream. The optical symbols stream shall then be delivered to the MDI, according to the transmit optical specifications (not included in this contribution). The lower optical power level shall correspond to pam_symbol = -1, and the higher optical power level to pam_symbol = +1.
 - pam_symbol = -1 \longleftrightarrow Lower optical power level (P_0)
 - pam_symbol = +1 \longleftrightarrow Higher optical power level (P_1)
 - $ER = P_1/P_0$
 - $OMA = P_1 - P_0$

Timing parameters



- The symbol transmission rate of the PHY shall be $1.0625 \times \text{bit-rate} \pm 0.01\%$
 - Bit-rate = 25 Gb/s: symbol-rate = 26.5625 GBd (= 170×156.25 MHz)
 - Bit-rate = 10 Gb/s: symbol-rate = 10.625 GBd (= 68×156.25 MHz)
 - Bit-rate = 5 Gb/s: symbol-rate = 5.3125 GBd (= 34×156.25 MHz)
 - Bit-rate = 2.5 Gb/s: symbol-rate = 2.65625 GBd (= 17×156.25 MHz)

Parameter	Value	Value	Value	Value
xMII Data-rate (Mb/s)	25000	10000	5000	2500
RS m	10	10	10	10
RS t	11	11	11	11
RS n	544	544	544	544
RS k	522	522	522	522
BaudRate MBd	26562.5	10625	5312.5	2656.25
Ref Clock (MHz)	156.25	156.25	156.25	156.25
Clock Factor	170	68	34	17
CWs/block	36	36	36	36
TX block length (symbols)	195840	195840	195840	195840
TX block duration (us)	7.37	18.43	36.86	73.73
FIFO latency (us)	0.45	1.13	2.25	4.51
PHD data-rate (Mb/s)	30.382	12.153	6.076	3.038
OAM net data-rate (Mb/s)	18.989	7.595	3.798	1.899

Physical header data (based on C/115)




- Physical Header Data (PHD): **same as C/115, where fields related with PDB offset and THP have been eliminated.**
- PHD is side information block embedded inside a transmit block that is used to exchange control and for negotiation of PCS and PMA parameters between two link partners
- All the PHD fields are transmitted from the least to the most significant bit and are transmitted from top to bottom of the table
- Each PHY has to deal with transmit and receive PHDs simultaneously. The prefix LOCPHD refers to the fields of the PHD to be included in the next transmit block transmitted to the link partner (from the local PHY). LOCPHD fields assigned by the state diagrams shall be sampled at the start of a transmit block to create the header included in that current transmit block
- The prefix REMPHD refers to the fields of the most recent PHD received, decoded, and validated with CRC from the link partner (from the remote PHY). The new values of REMPHD fields shall be available to the state diagrams and registers immediately after reception, decoding, and validation of the entire PHD and before the decoding of first RS codeword of the next received transmit block

Field name	# of bits	Description	Valid values
PHD.TX.NEXT.MODE	3	Transmission mode of the next Transmit Block, indicated to link partner to align its reception	0: normal transmission 1: BER test mode transmission 2 through 7: reserved
PHD.RX.LINKSTATUS	1	Indicates whether the local PHY is able to receive 65B data blocks with reliability. The value of this field is determined by the PHY quality monitor state diagram. The local PHY uses this received PHD field to determine the value of the variable rem_rcvr_status	0: NOT_OK 1: OK
PHD.RX.HDRSTATUS	1	Indicates whether the local PHY is able to receive the PHD from its link partner with reliability. The value of this field is determined by the local PHD reception monitor state diagram. The local PHY uses this received PHD field to determine the value of the variable rem_rcvr_hdr_lock	0: NOT_OK 1: OK
PHD.RX.LINKMARGIN	8	The value of this field is determined by the PHY quality monitor state diagram in response to link margin estimation as defined (TBD). Upon PHD reception, the field is stored in bits 3.XXX.7:0 (see e.g. see 45.2.3.55.1)	This field is fixed-point formatted (8, 3) and is provided in log2 units (see e.g. 115.3.7.2 and 115.3.8)
PHD.CAP.LPI	1	This field indicates the PHY supports and has enabled EEE, and that it is able to transmit and receive Low Power Idle	0: EEE is not supported or is disabled 1: EEE is supported and is enabled
PHD.CAP.OAM	1	This field indicates the PHY supports and has enabled OAM, and that it is able to transmit and receive management information by using the PHD.OAM.* fields	0: OAM is not supported or is disabled 1: OAM is supported and is enabled
	65	Reserved	0
PHD.OAM.DATA0	12	OAM message data field 0	0x000 through 0xFFFF
PHD.OAM.MSGT	1	OAM message identification bit	0 or 1
PHD.OAM.MERT	1	OAM STA read identification bit	0 or 1
PHD.OAM.PHYT	1	OAM PHY reception identification bit	0 or 1
	1	Reserved	0
PHD.OAM.DATA1	16	OAM message data field 1	0x0000 through 0xFFFF
PHD.OAM.DATA2	16	OAM message data field 2	0x0000 through 0xFFFF
PHD.OAM.DATA3	16	OAM message data field 3	0x0000 through 0xFFFF
PHD.OAM.DATA4	16	OAM message data field 4	0x0000 through 0xFFFF
PHD.OAM.DATA5	16	OAM message data field 5	0x0000 through 0xFFFF
PHD.OAM.DATA6	16	OAM message data field 6	0x0000 through 0xFFFF
PHD.OAM.DATA7	16	OAM message data field 7	0x0000 through 0xFFFF
PHD.OAM.DATA8	16	OAM message data field 8	0x0000 through 0xFFFF
Total	224		

CRC16 calculation (same as C/115)

- The 224-bit PHD is appended with 16 cyclic redundancy check bits (CRC16) to add error detection capability. The appended CRC16 shall be computed from the PHD bits and shall produce the same result as the implementation shown in the following figure. The generator polynomial is $(x + 1) \cdot (x^{15} + x + 1)$.
- The shift register elements S_0, \dots, S_{15} , shall be initialized with the value of 0x0000 for each PHD. The 224 PHD bits, in transmit bit order, are then used to compute the CRC16 with the multiplexer configured to CRCgen. After the 224 bits have been serially processed, the multiplexer is configured to CRCout and the 16 stored values are the CRC16. CRC16 is transmitted in order from S_{15} to S_0 .



Reed-Solomon code — specification, 1/3



- Reed-Solomon code operates over the Galois Field $GF(2^{10})$ where the symbol size is 10 bits. The encoder processes k message symbols to generate $2t$ parity symbols, which are then appended to the message to produce a codeword of $n = k + 2t$ symbols. RS code is denoted $RS(n, k)$, where $n = 544$, and $k = 522$. Therefore the correction capability of the code is $t = 11$ symbols.
- The code is based on the generating polynomial given by the following equation, where a is a primitive element of the finite field defined by the primitive polynomial $x^{10} + x^3 + 1$ (0x409).

$$g(x) = \prod_{j=1}^{2t} (x - \alpha^j) = \sum_{i=0}^{2t} g_i \cdot x^i$$

- The following equation defines the message polynomial $m(x)$ whose coefficients are the message symbols m_{k-1} to m_0 .

$$m(x) = \sum_{i=1}^k m_{k-i} \cdot x^{n-i}$$


- Each message symbol m_i is the bit vector $(m_{i,9}, m_{i,8}, \dots, m_{i,1}, m_{i,0})$, which is identified with the element of the finite field. The message symbols are composed of the bits of the aggregated 80 65B blocks with 20-bit PHD block, such that bit 0 of the 65B block 0 is the bit 0 of m_{k-1} and bit 19 of PHD block is the bit 9 of m_0 . The first symbol input to the encoder is m_{k-1} .
- The following equation defines the parity polynomial $p(x)$ whose coefficients are the parity symbols p_{2t-1} to p_0 .

$$p(x) = \sum_{i=1}^{2t} p_{2t-i} \cdot x^{2t-i}$$

Reed-Solomon code — specification, 2/3



- The parity polynomial is the remainder from the division of $m(x)$ by $g(x)$. This may be computed using the shift register implementation illustrated in the following figure.
- The outputs of the delay elements are initialized to zero prior to the computation of the parity for a given message. After the last message symbol, m_0 , is processed by the encoder, the outputs of the delay elements are the parity symbols for that message.
- The codeword polynomial $c(x)$ is then the sum of $m(x)$ and $p(x)$ where the coefficient of the highest power of x , $c_{n-1} = m_{k-1}$ is transmitted first and the coefficient of the lowest power of x , $c_0 = p_0$ is transmitted last. The first bit transmitted from each symbol is bit 0.



Reed-Solomon code — specification, 3/3




- The coefficients of the generator polynomial for the Reed-Solomon code are presented in the following table

Coefficients of the generator polynomial g_i (decimal)

i	g_i	i	g_i
0	807	12	374
1	280	13	544
2	944	14	374
3	621	15	482
4	3	16	555
5	177	17	976
6	365	18	452
7	657	19	899
8	813	20	783
9	1010	21	513
10	712	22	1
11	466		

Reed-Solomon code — performance

1 bits/dim, PAM-2, 0.959559 b/s/Hz/dim, RS(544,522) m=10, AWGN, Error-rate vs. SNR norm



Channel: AWGN
 RS(544,522) m = 10
 Spect. Eff.: 0.959559 b/s/Hz/dim
 Shannon gap (BER = 1e-12): 6.62 dB
 Capacity bound gap (BER = 1e-12): 3.76 dB
 SNR (BER = 1e-12): 11.07 dB
 Coding gain (BER = 1e-12): 5.55 dB
 Input SER (BER = 1e-12): 0.000175115

PMA state variables (based on C/115), 1/2



NOTE 1: In the following state variables definition, xMII means XGMII for 2.5, 5 and 10 Gb/s operation or 25GMII for 25 Gb/s operation

NOTE 2: New variables or changes in variables definition wrt C/115 will be indicated in **dark blue**

- **hdr_crc16_status:** result of the CRC16 evaluation for a received PHD from the link partner; this variable is assigned for each received PHD.

Values:

- OK: The most recent received PHD is correct as determined by CRC16 verification
- NOT_OK: The most recent received PHD is not correct as determined by CRC16 verification

- **hdr_fail_count:** variable used to count the reception of contiguous erroneous PHDs.

- **link_control:** variable that controls PMA functional operation. All state diagrams respond to the open-ended link_control = DISABLE.

Values:

- DISABLE: Prevent operation of PMA sublayer
- ENABLE: Permit operation of PMA sublayer

- **link_status:** variable that is set by the link monitor state diagram and used by PHY TX and PHY RX control state diagrams to control the 64B/65B encoder and 64B/65B decoder operation.

Values:

- OK: The link has been established between link partners with data reliability in both communication directions
- FAIL: Link is not established (one or both directions are not providing reliable payload data decoding)

- **loc_rcvr_hdr_lock:** variable set by the local PHD reception monitor state diagram to indicate the reliability of PHD reception.

Values:

- OK: Local PHD reception is reliable
- NOT_OK: Local PHD reception is unreliable

- **loc_rcvr_status:** variable set by the PHY quality monitor state diagram to indicate correct or incorrect **65B data blocks** decoding by the local PHY receiver.

Values:

- OK: **65B data blocks** reception by the local PHY is reliable
- NOT_OK: **65B data blocks** reception by the local PHY is unreliable

- **new_rxblock_event:** signal sent by the PHY receiver to indicate the start of reception of a new Transmit Block. This event persists only long enough to cause one state diagram transition.

- **new_rxphd_event:** signal sent by the PHY receiver to indicate the complete reception (including data decoding and CRC16 checking) of a new PHD from the link partner. This event persists only long enough to cause one state diagram transition.

- **pma_reset:** variable that causes reset of all PMA functions. PMA reset occurs with power on or the PHY reset being set to one (bit 1.0.15). All state diagrams respond to the open-ended pma_reset = ON.

Values:

- ON: Reset is asserted
- OFF: Reset is deasserted

PMA state variables (based on C/115), 2/2



- **rcvr_clock_lock**: variable set by the PMA receive clock recovery function to indicate that the clock has been properly recovered from the received signal.

Values:

- OK: Clock is stable and phase adjusted for sampling the received signal
- NOT_OK: Clock has not been recovered from the received signal and/or it is not stable

- **rcvr_hdr_lock**: variable set by the PHD monitor state diagram to indicate the reliability of both the PHD transmission from local to remote PHY and the PHD reception from remote to local PHY.

Values:

- OK: PHD transmission and reception are reliable
- NOT_OK: PHD transmission or reception are unreliable

- **rem_rcvr_hdr_lock**: variable set by the remote PHD reception monitor state diagram to indicate the reliability of PHD reception in the remote PHY (link partner).

Values:

- OK: Link partner PHD reception is reliable
- NOT_OK: Link partner PHD reception is unreliable

- **rem_rcvr_status**: variable set by the link monitor state diagram. It indicates the receiver status of the remote (link partner) PHY **65B blocks** decoding.

Values:

- OK: **65B blocks** reception by the remote PHY is reliable
- NOT_OK: **65B blocks** reception by the remote PHY is unreliable

- **rx_xmii_enable**: variable set by the PHY RX control state diagram to control the 64B/65B decoder operation.

Values:

- TRUE: The 64B/65B decoder receives **65B blocks** from the link partner and decodes them into **xMII** receive data stream transfers
- FALSE: The 64B/65B decoder does not decode received **65B blocks** from the link partner

- **sotxb_synch**: variable set by the PMA receive clock recovery function to indicate synchronization with the start of Transmit Blocks.

Values:

- OK: Synchronization **with the start of Transmit Blocks** has been achieved
- NOT_OK: Synchronization has not been achieved

- **tx_enable**: variable set by the PHY TX control state diagram to enable PCS and PMA transmission.

Values:

- TRUE: PCS and PMA transmission is enabled
- FALSE: PCS and PMA transmission is disabled

- **tx_xmii_enable**: variable set by the PHY TX control state diagram to control the 64B/65B encoder operation.

Values:


- TRUE: The 64B/65B encoder encodes the **xMII** transmit data stream transfers into **65B blocks**
- FALSE: The 64B/65B encoder does not encode the **xMII** transmit data stream. **Idles** are encoded in transmitted **65B blocks**

- **tx_xmii_idle**: variable that indicates the idle status of the **xMII** transmit data path.


Values:

- TRUE: **Idle** is encoded in the **xMII** transmit stream
- FALSE: **Idle** is not encoded in the **xMII** transmit stream

PMA state diagrams (based on C/115), 1/3




PHY TX control state diagram
(only variable names are changed)




PHY RX control state diagram
(variable names changed, eliminated PMARX_THP_INIT state because THP is not used / initialized)

PMA state diagrams (based on C/115), 2/3




Link monitor state diagram




Local PHD reception monitor state diagram
 (variable name changed, transition from LOCHDR_LOCK to LOCHDR_UNLOCK is not triggered by start of transmit block synch fail)

PMA state diagrams (based on C/115), 3/3



Remote PHD reception monitor state diagram



PHD monitor state diagram

PHY receiver quality assessment (based on C/115)




- PHY quality criterion: 65B data blocks reception is reliable when the RS frame error ratio (RSFER) is less than 5×10^{-10} after RS decoding
- The assessment of PHY quality criterion may be based on estimation of the noise variance at the PAM2 decoder decision points $E[n_d^2]$, which expressed in base-2 logarithmic units has to be lower than a given threshold T_{LM} to have reliable local reception of 65B data blocks. If the condition $\log_2(E[n_d^2]) < T_{LM}$ holds, the variable `loc_rcvr_status` is assigned the value OK. The noise variance at the PAM2 decoder can be estimated either by measuring the Modulation Error Ratio (MER) at the decision points or measuring the ratio of corrected symbols per codeword carried out by the RS decoder. The value of the threshold and the information used to estimate the RS decoder noise variance is implementation dependent.
- The local PHY reports the link margin to the link partner by using the field `LOCPHD.RX.LINKMARGIN`. The link margin (LM) is defined as the SNR margin relative to the SNR required for reception of RS-coded PAM2 with the indicated quality criterion.
- Following equation formally defines the link margin:

$$LM = T_{LM} - \log_2(E[n_d^2])$$

PHY quality monitor state diagram (based on C/115)



- **new_link_margin_event**: signal sent by the PHY receiver to indicate a new estimation of detector link margin is available. This event persists only long enough to cause one state diagram transition.
 - **link_margin**: variable set by the PHY receiver containing the value of the last link margin estimation.
- Values:
- Any value determined per Link Margin equation previously defined



PHY quality monitor state diagram
 (PMAMON_DISABLE and PMAMON_WAITING states have been merged because PMA check quality does not need wait until THP is initialized because THP is not used / initialized)

OAM channel (same as C/115)



- Physical layer Operations Administration Maintenance (OAM) channel has been included as capability (optional implementation) in several automotive PHY standards: 802.3bv, 802.3bp, 802.ch
- Physical layer OAM channel provides a mechanism to reliably exchange messages between station management entity (STA) peers attached to link partners. The OAM message exchange occurs in the PCS, as part of the PHD, and does not impact the normal xMII to xMII data transmission. Moreover, the OAM message exchange is not affected by EEE operation
- The OAM channel utilizes OAM transmit and receive registers accessible via the MDIO
- The presented PHD structure and transmit block structure are 100% compatible with the already specified 1000BASE-H OAM channel in C/115.9 and MDIO registers in C/45.2.3.50 and C/45.2.3.51
- OAM channel as specified 802.3bv is proposed to be reused with NO changes for OMEGA PHYs

OAM channel — registers (same as C/115)



- Following tables are the existing OAM registers definition for 1000BASE-H.
- To be decided by the TF if new set of registers have to be redefined with the same content or the existing registers can be reused (with editorial changes) by any physical layer that implements the same OAM channel, i.e. 1000BASE-H, OMEGA, etc.
- In the following state variables definition, it is assumed the register addresses of 1000BASE-H OAM. Different addresses should be used in case of separate registers.

Table 45–220—1000BASE-H OAM transmit register bit definitions

Bit(s)	Name	Description	R/W ^a
3.500.15	TXO_REQ	1 = Transmission of a 1000BASE-H OAM message is pending; write as one to request transmission 0 = 1000BASE-H transmit registers are available for a new 1000BASE-H OAM message	R/W, SC
3.500.14	TXO_PHYT	The identifier of the last 1000BASE-H OAM message received by the remote PHY	RO
3.500.13	TXO_MERT	The identifier of the last 1000BASE-H OAM message read by the STA attached to the remote PHY	RO
3.500.12	TXO_MSGT	1000BASE-H OAM message identifier; its value change with each new transmitted message	RO
3.500.11:0	TXO_DATA0	Transmit 1000BASE-H OAM message first 12 bits	R/W
3.501.15:0	TXO_DATA1	Transmit 1000BASE-H OAM message 16-bit data word 1	R/W
3.502.15:0	TXO_DATA2	Transmit 1000BASE-H OAM message 16-bit data word 2	R/W
3.503.15:0	TXO_DATA3	Transmit 1000BASE-H OAM message 16-bit data word 3	R/W
3.504.15:0	TXO_DATA4	Transmit 1000BASE-H OAM message 16-bit data word 4	R/W
3.505.15:0	TXO_DATA5	Transmit 1000BASE-H OAM message 16-bit data word 5	R/W
3.506.15:0	TXO_DATA6	Transmit 1000BASE-H OAM message 16-bit data word 6	R/W
3.507.15:0	TXO_DATA7	Transmit 1000BASE-H OAM message 16-bit data word 7	R/W
3.508.15:0	TXO_DATA8	Transmit 1000BASE-H OAM message 16-bit data word 8	R/W

^aR/W = Read/Write, RO = Read only, SC = Self-clearing

Table 45–221—1000BASE-H OAM receive register bit definitions

Bit(s)	Name	Description	R/W ^a
3.509.15	RXO_VAL	1 = A new 1000BASE-H OAM message has arrived 0 = No new 1000BASE-H OAM message arrived since either last message was read or PMA reset	RO
3.509.14:13	Reserved	Value always 0	RO
3.509.12	RXO_MSGT	Identifier of the received 1000BASE-H OAM message	RO
3.509.11:0	RXO_DATA0	Receive 1000BASE-H OAM message first 12 bits	RO
3.510.15:0	RXO_DATA1	Receive 1000BASE-H OAM message 16-bit data word 1	RO
3.511.15:0	RXO_DATA2	Receive 1000BASE-H OAM message 16-bit data word 2	RO
3.512.15:0	RXO_DATA3	Receive 1000BASE-H OAM message 16-bit data word 3	RO
3.513.15:0	RXO_DATA4	Receive 1000BASE-H OAM message 16-bit data word 4	RO
3.514.15:0	RXO_DATA5	Receive 1000BASE-H OAM message 16-bit data word 5	RO
3.515.15:0	RXO_DATA6	Receive 1000BASE-H OAM message 16-bit data word 6	RO
3.516.15:0	RXO_DATA7	Receive 1000BASE-H OAM message 16-bit data word 7	RO
3.517.15:0	RXO_DATA8	Receive 1000BASE-H OAM message 16-bit data word 8	RO

^aRO = Read only

OAM channel – state variables (same as C/115)



- **oam_cap**: enables OAM channel functionality. This variable is set to TRUE when PHD.CAP.OAM of both transmit and receive PHD is TRUE. Otherwise, it is FALSE.
Values:
 - TRUE: Both local and remote PHY have OAM ability and OAM functionality is enabled in both PHYs
 - FALSE: Either local or remote PHY does not have OAM ability or it is disabled
- **read_RXOAM_DATA8_event**: event indicating that the OAM receive register 3.517 has been read. This event persists only long enough to cause one state diagram transition.
- **rxphd_mert**: variable equivalent to the field REMPHD.OAM.MERT.
Values: 0 and 1
- **rxphd_msgt**: variable equivalent to the field REMPHD.OAM.MSGT. It contains the message identification bit of the OAM message.
Values: 0 and 1
- **rxphd_oamudat**: variable equivalent to the fields REMPHD.OAM.DATAx. It is the payload of the OAM message.
Values: Any value
- **rxphd_phyt**: variable equivalent to the field REMPHD.OAM.PHYT. Values: 0 and 1
- **rxr_msgt**: variable equivalent to the bit RXO_MSGT of register 3.509. This bit is the message identification bit of the message contained in the OAM receive registers.
Values: 0 and 1
- **rxr_oamudat**: variable equivalent to the RXO_DATA0 field of register 3.509 concatenated with the RXO_DATA1 through RXO_DATA8 of OAM receive registers 3.510 through 3.517.
Values: Any value
- **rxr_rxval**: variable equivalent to the bit RXO_VAL of register 3.509 that indicates the presence of a valid message in the OAM receive registers.
Values:
 - 1: There is a valid message in the receive registers, which is pending for processing by the local STA
 - 0: There is no valid message in the receive registers
- **txphd_mert**: variable equivalent to the field LOCPHD.OAM.MERT. It informs the remote PHY of the message identification bit of the last message that has been received by the local STA.
Values: 0 and 1
- **txphd_msgt**: identifier of the message being transmitted by the local PHY. This variable is equivalent to the field LOCPHD.OAM.MSGT.
Values: 0 and 1
- **txphd_oamudat**: variable equivalent to the content of the field LOCPHD.OAM.DATAx.
Values: Any value
- **txphd_phyt**: variable equivalent to the field LOCPHD.OAM.PHYT. It indicates to the remote PHY the message identification bit of the last valid message written to the receive registers.
Values: 0 and 1
- **txr_mert**: variable equivalent to the bit TXO_MERT of register 3.500. This bit is updated in OAMTX_TRANSMIT state to reflect the identifier of the last message that has been acknowledged by the remote STA.
Values: It alternates between values 0 and 1
- **txr_msgt**: variable equivalent to the bit TXO_MSGT of register 3.500. This bit is updated in OAMTX_TRANSMIT state to reflect the message identification bit of the message being sent by the local PHY.
Values: It alternates between values 0 and 1
- **txr_oamudat**: variable equivalent to the TXO_DATA0 field of register 3.500 concatenated with the TXO_DATA1 through TXO_DATA8 of OAM transmit registers 3.501 through 3.508.
Values: Any value
- **txr_phyt**: variable equivalent to the bit TXO_PHYT of register 3.500. This bit is updated in OAMTX_TRANSMIT state to reflect the identifier of the last message that has been acknowledged by the remote PHY.
Values: It alternates between values 0 and 1
- **txr_txreq**: variable equivalent to the bit TXO_REQ of register 3.500. It indicates if the local STA is requesting the transmission of a new OAM message.
Values: 0: There is no message transmission request

OAM channel — states diagrams (same as C/115)




Figure 115-42—PHY 1000BASE-H OAM transmit control state diagram




Figure 115-43—PHY 1000BASE-H OAM receive control state diagram

OAM channel — robust protocol, A to B



- Step 1: STA A writes OAM message TX registers and raises TXO_REQ bit

OAM Tx registers


TXO_REQ	TXO_PHYT	TXO_MERT	TXO_MSGT
1	<i>a</i>	<i>a</i>	<i>a</i>

PHY A

OAM Rx registers

RXO_VAL	RXO_MSGT
0	<i>a</i>

PHY B



Extracted from contribution https://www.ieee802.org/3/bv/public/Mar_2015/perezaranda_3bv_4_0315.pdf

OAM channel — robust protocol, A to B



- Step 2: local PHY (A) acknowledges the message to local STA (A) with TXO_MSGT and transmit it to the remote PHY (B)

OAM Tx registers


TXO_REQ	TXO_PHYT	TXO_MERT	TXO_MSGT
0	<i>a</i>	<i>a</i>	<i>!a</i>

PHY A

OAM Rx registers

RXO_VAL	RXO_MSGT
0	<i>a</i>

PHY B



Extracted from contribution https://www.ieee802.org/3/bv/public/Mar_2015/perezaranda_3bv_4_0315.pdf

OAM channel — robust protocol, A to B



- Step 3: remote PHY (B) receives the message. It signals the reception of the message to the remote STA (B) using RXO_VAL and sends PHYT bit to the local PHY (A) indicating to local STA (A) messages was received by remote PHY (B)

OAM Tx registers


TXO_REQ	TXO_PHYT	TXO_MERT	TXO_MSGT
0	<i>!a</i>	<i>a</i>	<i>!a</i>

OAM Rx registers

RXO_VAL	RXO_MSGT
1	<i>a</i>

PHY A

PHY B



Extracted from contribution https://www.ieee802.org/3/bv/public/Mar_2015/perezaranda_3bv_4_0315.pdf

OAM channel – robust protocol, A to B



- Step 4: remote STA (B) reads the message, the message read toggle bit MERT reaches the local PHY (A) and it is signaled through the local OAM TX registers to local STA (A) indicating the remote STA (B) already read the message

OAM Tx registers


TXO_REQ	TXO_PHYT	TXO_MERT	TXO_MSGT
0	<i>!a</i>	<i>!a</i>	<i>!a</i>

OAM Rx registers

RXO_VAL	RXO_MSGT
0	<i>!a</i>


PHY A

PHY B



Extracted from contribution https://www.ieee802.org/3/bv/public/Mar_2015/perezaranda_3bv_4_0315.pdf

OAM channel – bidirectional TX with handshaking



Extracted from contribution https://www.ieee802.org/3/bv/public/Mar_2015/perezaranda_3bv_4_0315.pdf

References



- [1] R. Pérez-Aranda, “OMEGA link model,” October 2020, [Online], Available: https://www.ieee802.org/3/cz/public/13_oct_2020/perezaranda_3cz_01_131020_link_model.pdf
- [2] R. Pérez-Aranda, “OMEGA 25 Gb/s link budget analysis,” October 2020, [Online], Available: https://www.ieee802.org/3/cz/public/27_oct_2020/perezaranda_3cz_01_271020_25G_link_budget.pdf
- [3] R. Pérez-Aranda, “OMEGA 10 Gb/s link budget analysis,” October 2020, [Online], Available: https://www.ieee802.org/3/cz/public/27_oct_2020/perezaranda_3cz_02_271020_10G_link_budget.pdf
- [4] R. Pérez-Aranda, “Impact of longer wavelengths fiber response in the 25 Gb/s link budget,” October 2020, [Online], Available: https://www.ieee802.org/3/cz/public/27_oct_2020/perezaranda_3cz_03_271020_25G_emb_impact.pdf



Thank you!