

100GBASE-KP4 Training Idea

Kent Lusted, Intel
Adee Ran, Intel
Matt Brown, AppliedMicro

(Regarding Comment #125)

Supporters

- Andre Szczepanek, Inphi
- Arash Farhood, Cortina Systems
- Brad Booth, Dell
- Dave Chalupsky, Intel
- Rich Mellitz, Intel
- Vasu Parthasarathy, Broadcom
- Ilango Ganga, Intel

Assumptions

- Based on P802.3bj Draft 1.0 and P802.3bh Draft 3.1
- Maximize compatibility with existing training mechanism in Clause 72.6.10
- Leverage the tight integration and inter-dependence of the PMA and PMD functions, specific to Clause 94.
 - Similar to PCS/PMA dependency in 10GBASE-T link training (Clause 55)

Leverage Existing

- Take 72.6.10 as the baseline
- Modify as given in the following slides

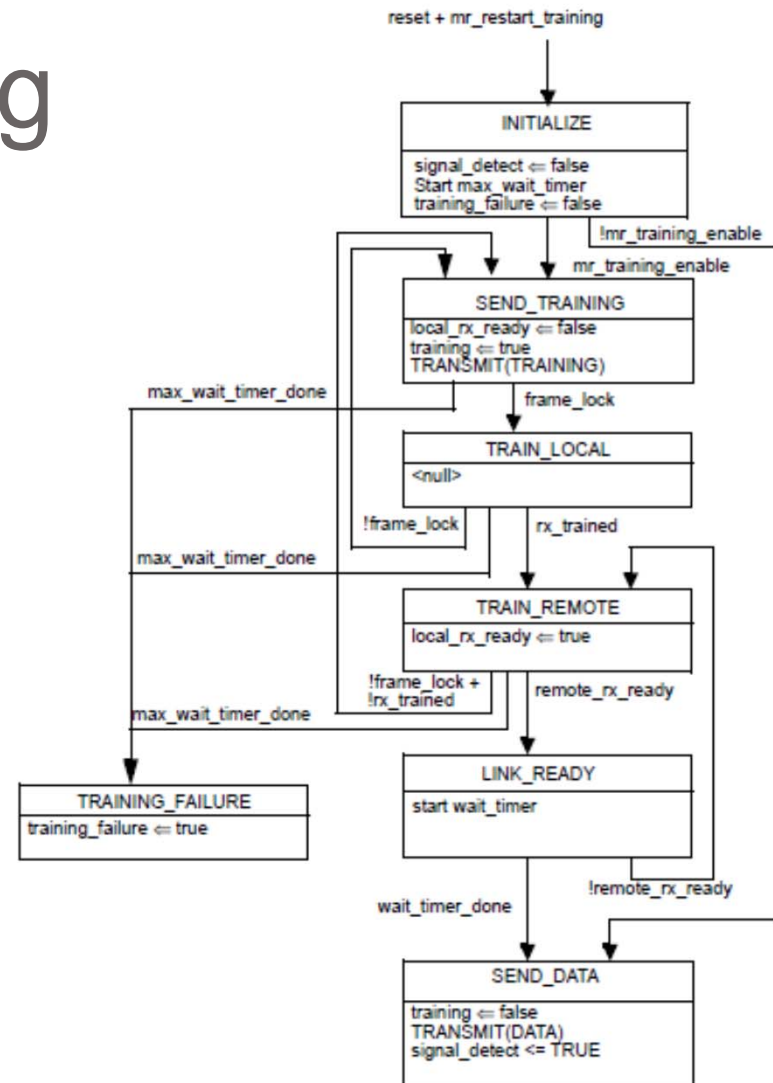


Figure 72-5—Training state diagram

Training Frame Structure

- Keep same basic format as Clause 72
 - Frame marker
 - Control channel
 - Coef update field
 - Status report field
 - Training pattern





Figure 72-2—Training frame structure

100GBASE-KP4 Change Preface

- Pack the Frame Marker, Control Channel and Training Pattern into 46 UI training frame words (TFW)
 - Each TFW corresponds to two full 46-bit terminated blocks (TB46).
 - Simplifies design (i.e. no gearbox)
 - Enables early data alignment during training period
 - Enables regular data recovery and fast switching to data mode
- Use PAM4 signaling in the training pattern
- Adapt status report field to expedite PMA alignment on data mode transition
- Add parity check to both control channel fields



Rapid Transition Of Training -> Data Mode

- Training Frame alignment to the Overhead Frame logic is important for rapid transition from the training mode to data mode
 - 100GBASE-KP4 PMA frame consists of 696 TB46
 - Proposed 100GBASE-KP4 training frame is 384 TB46
 - $\text{Gcd}(696, 384) = 24$
 - The offset is periodical with a cycle of $696/24=29$ training frames
- Training frames shall always be transmitted with offset which is $24 \cdot k$, $k \in \{0 \dots 28\}$ words relative to the 40-bit overhead
 - Other offsets are illegal
 - Assuming the first training frame starts at offset 0 ($k=0$), then the second training frame starts at offset 384 ($k=16$), and subsequent training frames start at offsets with k equal to 3, 19, 6, 22 ...

Frame Marker and Control Channel

- To simplify link training in the link establishment process when TRANSMIT(TRAINING), the PAM4 multi-level signaling shall not be used for the Frame Marker, Coefficient Update, and Status Report fields.
 - Make it full-swing, i.e. “NRZ” like
- During the transmission of the Frame Marker, Coefficient Update, and Status Report fields, the PAM4 transmitter *shall bypass the overhead frame, termination block, gray coding, and $1/(1+D)$ mod 4 precoding stages of the PMA transmit and receive functional specifications.*
 - *Therefore, the output levels shall be restricted to level -1 for a 0 and level +1 for a 1 to enable easy receiver lock to the training pattern over poor quality and non-equalized channels.*

Frame Marker – TFW #1

- Frames are delimited by the 46 UI pattern, hexadecimal 3FFFFFFF800000 (ones transmitted first), as expressed in 13.59375 Gbd symbols.
- This pattern does not appear in the control channel or the training pattern and therefore serves as a unique indicator of the start of a training frame.

training frame word	45:26	35:26	25:16	15:6	5:0		
1	10'b1	10'b1	3'b1 7'b0	10'b0	6'b0		frame marker

Control Channel Cells

- Signaled using differential Manchester encoding (DME) like in Clause 72.6.10.2.2
- Each DME symbol contains 2 DME transition positions.
 - Each transition position is the mid-point of the cell.
- The data cell length is 10 100GBASE-KP4 UI. (~736ps)
 - Approximately the same duration at 10GBASE-KR cells
 - Transition position is 5 UI

Control Channel Encoding

- Control channel uses 10 TFWs
 - TFW #2 to #11
- Pack 4 cells of control channel into 1 TFW
 - Cells take definitions and ordering as shown in Table 72-4 and Table 72-5
 - Coef Update Field before Status Update Field
 - Highest # cell first, in descending order
 - Uses 40 UI of 46 UI in 1 TFW

What about the Remaining 6 UI?

- Define the last 6 UI in each of TFW #2-11 as overhead cell
 - Transition position is 3 UI
- Set overhead cell to a DME logic '1' of 6 UI width
 - Preserve the DC balance on the line
 - 000111 or 111000, depending on previous cell value

Coef Update Field

Table 72-4—Coefficient update field

Cell(s)	Name	Description
15:14	Reserved	Transmitted as 0, ignored on reception.
13	Preset	1 = Preset coefficients 0 = Normal operation
12	Initialize	1 = Initialize coefficients 0 = Normal operation
11:8	Reserved	Transmitted as 0, ignored on reception.
5:4	Coefficient (+1) update	<u>5</u> <u>4</u> 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold
3:2	Coefficient (0) update	<u>3</u> <u>2</u> 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold
1:0	Coefficient (-1) update	<u>1</u> <u>0</u> 1 1 = reserved 0 1 = increment 1 0 = decrement 0 0 = hold

11:7 Reserved
6 Parity Check

Status Report Field

- 16 -> 20 cells
 - 5 TFWs
 - 10 TB46
- Add new features
- Keep coef status and Receiver Ready
- Cell ordering not finalized

Cell(s)	Name	Description
19	Parity Check	Parity calculation for Status Report Field
18:14	EEE State	Current EEE state of local transmitter, if EEE is implemented.
13:12	Training Frame Countdown	Number of training frames remaining before link training process transitions to data mode
11:7	PMA Alignment Offset	Relative location of the next training frame within the PMA frame
6	Receiver ready	1 = The local receiver has determined that training is complete and is prepared to receive data. 0 = The local receiver is requesting that training continue.
5:4	coefficient (+1) status	<u>5</u> <u>4</u> 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated
3:2	coefficient (0) status	<u>3</u> <u>2</u> 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated
1:0	Coefficient (-1) status	<u>1</u> <u>0</u> 1 1 = maximum 1 0 = minimum 0 1 = updated 0 0 = not_updated

Two New Parity Check Fields

- This is an improvement over the original clause 72 rules
 - Increases protection against false acceptance of sensitive messages, e.g. preset, init, receiver ready
 - Guarantees DC balance of DME cells during training
- Use cell 6 of the coef update field and cell 19 of status report field to encode a parity check for each respective field
 - The two parity bits are calculated for each field (coef update, status report) separately, such that each field *as a whole, including the parity bit* has even parity
 - Parity is the number of logical-one cells in the field, modulo 2 (not including the overhead bits)
- If a parity violation is detected within the bounds of the respective field in a given training frame, the contents of the control channel for that frame shall be ignored.


New Status Report Cells

- **EEE State (Cells 18:14)**
 - Current EEE state of local transmitter, if EEE is implemented
 - Otherwise, reserved and set to 0.
 - See brown_01_0712.pdf for more details
- **Training frame countdown counter (Cells 13:12)**
 - Used to signal the transition from training to data mode.
 - Start at 3, decrement toward 0 during the last 3 frames sent
 - When a frame is sent with this value = 0, after the last TFW of training pattern (#102) is sent, transmission switches to the PMA frame

New Status Report Cells (2)

- PMA Alignment Offset -- PAO (Cells 11:6)
 - Represents the relative location measured in TB46 (modulus 696) of the next training frame within the PMA frame
 - PAO encodes the relative location as a 5-bit integer
 - The start of the next training frame is $24 * \text{PAO}$
 - 0: marker aligned with 40-bit overhead
 - 1: marker is at offset of $24*1=24$ termination blocks from 40-bit overhead
 - ...
 - 28: marker is at offset $24*28=672$ termination blocks from 40-bit overhead
 - 29 to 31: invalid, never transmitted, ignored on reception
 - To shift instantly to data mode after the last training frame ends, RX needs to know the relative offset of the next TB46 from the 40-bit overhead within the PMA frame

PMA Alignment Offset Example



Frame Marker and Control Channel Mapping

Symbol TX order →

Block TX order ↓

TFW	T-block	45:26	35:26	25:16	15:6	5:0	
0	0:1	10'b1	10'b1	3'b1 7'b0	10'b0	6'b0	frame marker
1	2:3	cell 15	cell 14	cell 13	cell 12	overhead	coef update
2	4:5	cell 11	cell 10	cell 9	cell 8	overhead	coef update
3	6:7	cell 7	cell 6	cell 5	cell 4	overhead	coef update
4	8:9	cell 3	cell 2	cell 1	cell 0	overhead	coef update
5	10:11	cell 19	cell 18	cell 17	cell 16	overhead	status report
6	12:13	cell 15	cell 14	cell 13	cell 12	overhead	status report
7	14:15	cell 11	cell 10	cell 9	cell 8	overhead	status report
8	16:17	cell 7	cell 6	cell 5	cell 4	overhead	status report
9	18:19	cell 3	cell 2	cell 1	cell 0	overhead	status report

Training Pattern Motivation

- Use the PMA transmit and receive functional specifications as currently defined in P802.3bj Draft 1.0 to enable the transmitter and receiver to exercise termination block, gray coding, and $1/(1+D) \bmod 4$ precoding stages.
 - Overhead framer is not exercised, but alignment is tracked through PMA Alignment Offset (PAO) cell
- Generate multi-level PAM4 signaling for receiver calibration
- Choose a pattern that is PMA termination block friendly

Training Pattern Details

- Apply termination, Gray coding, $1/(1+D)$ mod 4 precoding and PAM4 mapping
- PRBS13 is used for training pattern generation, followed by its inverse
 - $G(x) = 1 + x^2 + x^{11} + x^{12} + x^{13}$
 - Full length is $2 \cdot (2^{13} - 1) = 16382$
 - Last 2 bits of the inverse PRBS13 are discarded, so PRBS uses $16380/45 = 364$ TB46
 - 4 initial seeds selected to create desired properties
- Pack each 45 bits of PRBS12 into TB46
- Training frame (frame marker, control channel and training pattern) contains 384 TB46

Training Pattern Initial State

- Seeds (16 bits, initial transmitted bits, LSB first)
 - PMD lane 0: 0x836F
 - PMD lane 1: 0x4007
 - PMD lane 2: 0xB974
 - PMD lane 3: 0xD3D4

Note: for the underlined characters, only the LSB is part of the seed; the 3 MSBs are calculated by the polynomial
- Using these four states creates a distinct DC balanced sequence for each lane
 - Helps with lane order identification later on.
 - Generate the full PRBS13 starting from the initial seed, and then another full PRBS13 from the same seed, inverted
 - PRBS state and precoder state shall not advance during frame marker and control channel transmission and is reset for each training frame.

Additional Training Pattern Details

Training Pattern Seed Mapping

- $G(x) = 1 + x^2 + x^{11} + x^{12} + x^{13}$

SEED:

S0=0x836F

TFW #10

Initial 90 bits PRBS13
(from L to R)

PMD Lane 0

11110 11011 00000 11110 01011 00110 01011 11111 11101_01101 10111 01111 00111 00111 00111 00100 01010 01000 11101_
2213200123113131122223313212123120131011030233
 201200011232103233330321102330113210010031123

After Gray coding

After precoding

S1=0x4007

PMD Lane 1

11100 00000 00001 01101 10110 00010 10001 00001 10101_01011 11101 10110 00001 11100 10110 01100 10111 11111_
2300000321320033030133311221320012311313112223
 2131313023022212213212323332113102101230102023

S2=0xB974

PMD Lane 2

00101 11010 01110 10101 11110 00000 01101 00000 00100_00110 11000 11001 01001 11111 10111 11011 01010 11000_
0323312111220001330003002130203312221221333200
 0330320101113132122221002300221233332023030220

S3=0xD3D4


PMD Lane 3

00101 01111 00101 10101 11110 10000 01101 01110 11111_01110 00110 10101 00010 00000 00111 00000 00010 10000_
0332203211221001332322312013333030001200003300
 0302003323332223033020320012121300001131312130

Notes:

- For the underlined characters, only the LSB is part of the seed; the 3 MSBs are calculated by the LFSR equation
- Precoder starts from state 0

Training Frame Example

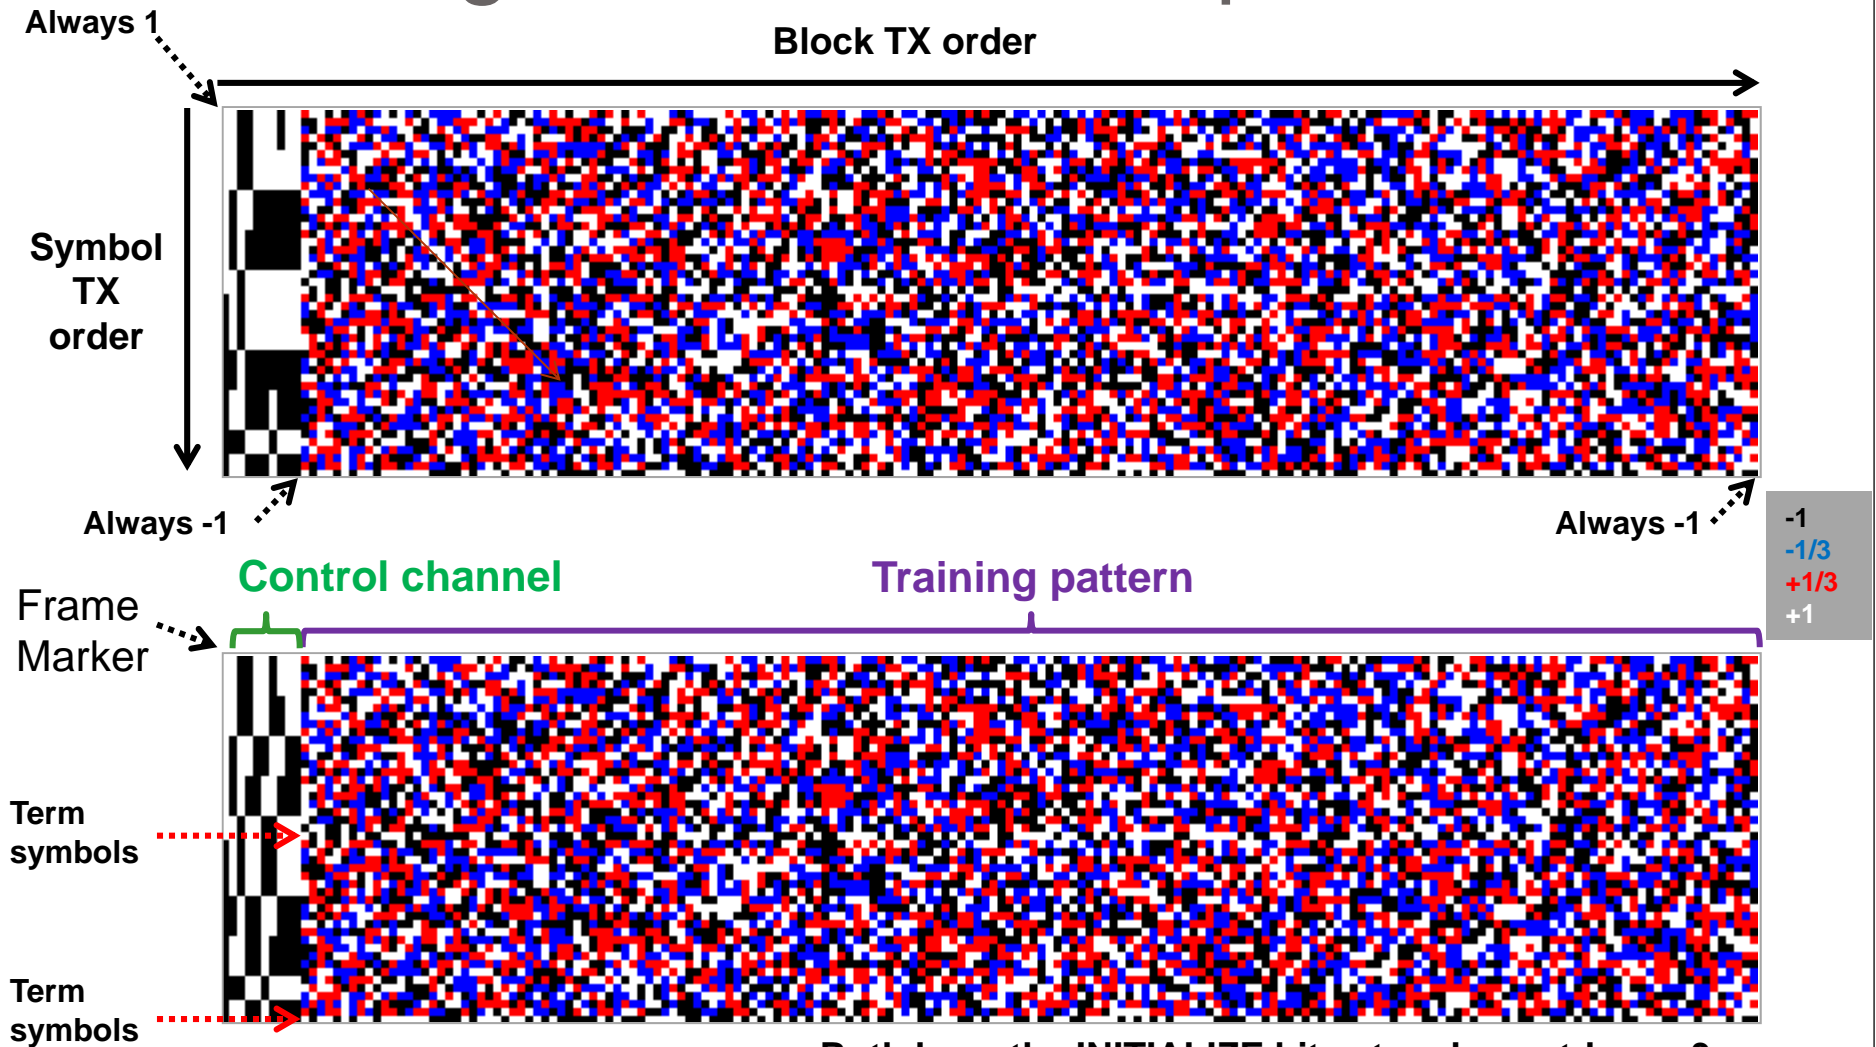


- Blue = no channel
- Red = after IL = $\sim 20\text{dB}$ @ 7GHz channel

Channel used "Thru_S14-P23-04-AB_S06-P20-10-CD_NNN.s4p" from

http://www.ieee802.org/3/100GCU/public/ChannelData/emerson_11_0928/meier_01_1011.pdf

Training Frame @ Un-Equalized RX



- Both have the INITIALIZE bit set and countdown=3.
- For the first frame, PAO=0, and for the second PAO=16.
- All other logical bits in the control channel are 0, so parity is 0 for the first frame and 1 for the second.
- Training pattern of lane 0 is used.

Training Frame Time Breakdown

- 1 Training frame contains 384 TB46
 - 8832 UI or 649.7 ns;
 - 95% of the duration is rich-content signal
- For reference, 1 training frame for 10GBASE-KR is ~425 nsec
- Link_fail_inhibit_timer (KR) = 500-510 msec
 - ~770k frames in 100GBASE-KP4
 - ~1.2E6 frames in 10GBASE-KR

Field	# of TB46	# of TFWs	TB46 start index	TB46 end index	TFW start index	TFW end index
Marker	2	1	0	1	0	0
coefficient request	8	4	2	9	1	4
status update	10	5	10	19	5	9
PRBS13+PRBS13i, truncated to 16380 bits	364	182	20	383	10	191

PMA frame length [TB46]	TF length [TB46]	GCD	Possible offsets	PAO width	PRBS length [TB46]	PRBS % of TF	TF length(ns)
696	384	24	29	5	364	94.8%	649.7

Conclusion

- 100GBASE-KP4 Training frame as proposed:
 - Reuses most of the existing 10GBASE-KR training mechanism
 - Leverages tight-integration of PMA and PMD architecture unique to 100GBASE-KP4
 - Provides full swing and PAM4 signaling
 - Supplies DC-balanced, lane-specific seeded training pattern
 - Offers parity check for coef update and status report fields
 - Enables fast and efficient transition to data mode
 - Facilities EEE signaling

BACKUP

PAM4 levels

- **PAM 4 consists of 4 levels mapped as follows:**
 - **0 maps to -1**
 - **1 maps to -1/3**
 - **2 maps to +1/3**
 - **3 maps to +1**

10GBASE-KR Link Establishment in Brief

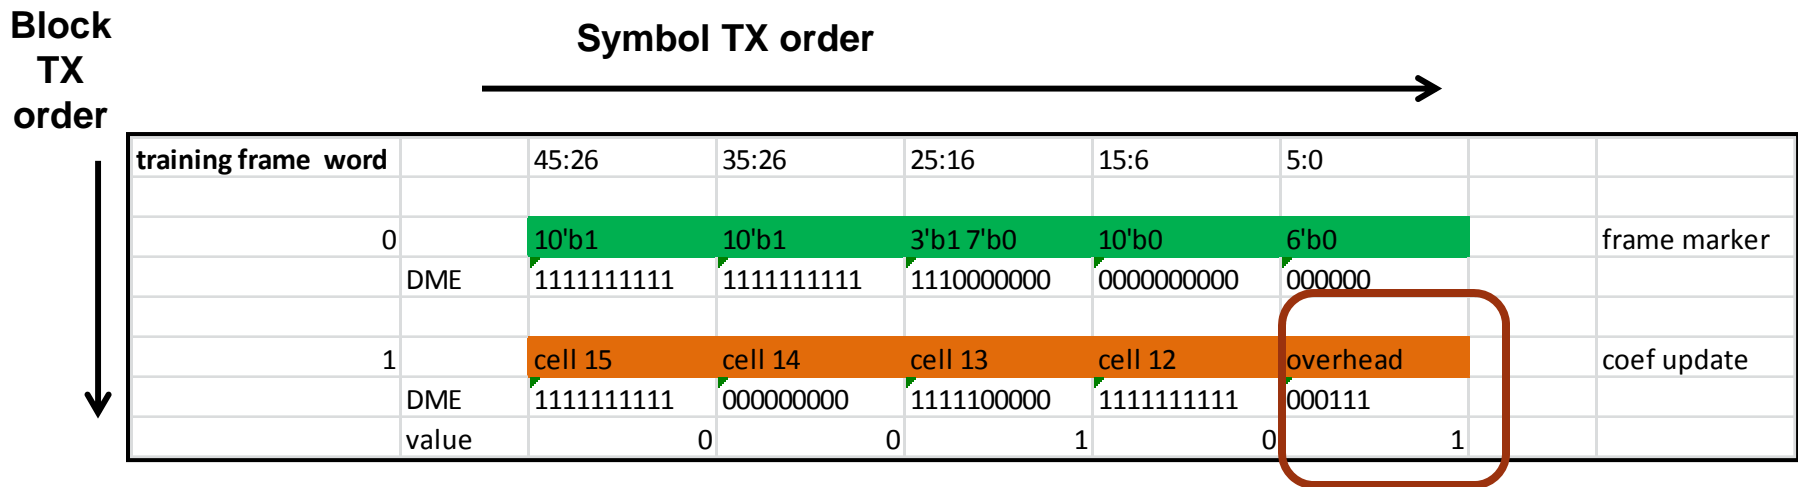
- The link is established using the following sequence:
 - (1) Auto-negotiate capabilities to link partner
 - (2) Send out training sequence to tune PHY for the channel's characteristics
 - Obtain Frame Lock
 - TX FFE handshake: Adapt tx coefficients to channel characteristics
 - DSP converged: Train RX to channel
 - Status Exchange: Ready or not?
 - (3) Go to data mode and Send out idle symbols

Control Channel Encoding Rules

- The control channel shall be transmitted using differential Manchester encoding (DME). The rules of differential Manchester encoding are as follows:
 - a) A data transition shall occur at each cell boundary.
 - b) A mid-cell data transition shall be used to signal a logical one.
 - c) The absence of a mid-cell data transition shall be used to signal a logical zero.
- If a coding violation is detected within the bounds of the control channel in a given training frame, the contents of the control channel for that frame shall be ignored.

TFW Overhead Cell Example 1

- Example 1: TFW#2 (Coef Update Field with PRESET asserted)
 - Cell 15 = 0, Cell 14 = 0, Cell 13 = 1, Cell 12 = 0



TFW Example 2

- Example 1: TFW#2 (Coef Update Field with nothing asserted)
 - Cell 15 = 0, Cell 14 = 0, Cell 13 = 0, Cell 12 = 0

Block TX order
↓

Symbol TX order →

training frame word	45:26	35:26	25:16	15:6	5:0	
0	10'b1	10'b1	3'b1 7'b0	10'b0	6'b0	frame marker
DME	1111111111	1111111111	1110000000	0000000000	000000	
1	cell 15	cell 14	cell 13	cell 12	overhead	coef update
DME	1111111111	0000000000	1111111111	0000000000	111000	
value		0	0	0	0	1

Parity Example 1

Symbol TX order



Block TX order



training frame word	45:26	35:26	25:16	15:6	5:0	
0	10'b1	10'b1	3'b1 7'b0	10'b0	6'b0	frame marker
DME	1111111111	1111111111	1110000000	0000000000	000000	
1	cell 15	cell 14	cell 13	cell 12	overhead	coef update
DME	1111111111	0000000000	1111111111	0000000000	111000	
value		0	0	0	0	1
2	cell 11	cell 10	cell 9	cell 8	overhead	coef update
DME	1111111111	0000000000	1111111111	0000000000	111000	
value		0	0	0	0	1
3	cell 7	cell 6	cell 5	cell 4	overhead	coef update
DME	1111111111	0000000000	1111111111	0000000000	111000	
value		0	0	0	0	1
4	cell 3	cell 2	cell 1	cell 0	overhead	coef update
DME	1111111111	0000000000	1111111111	0000000000	111000	
value		0	0	0	0	1

Parity Example 2

Symbol TX order

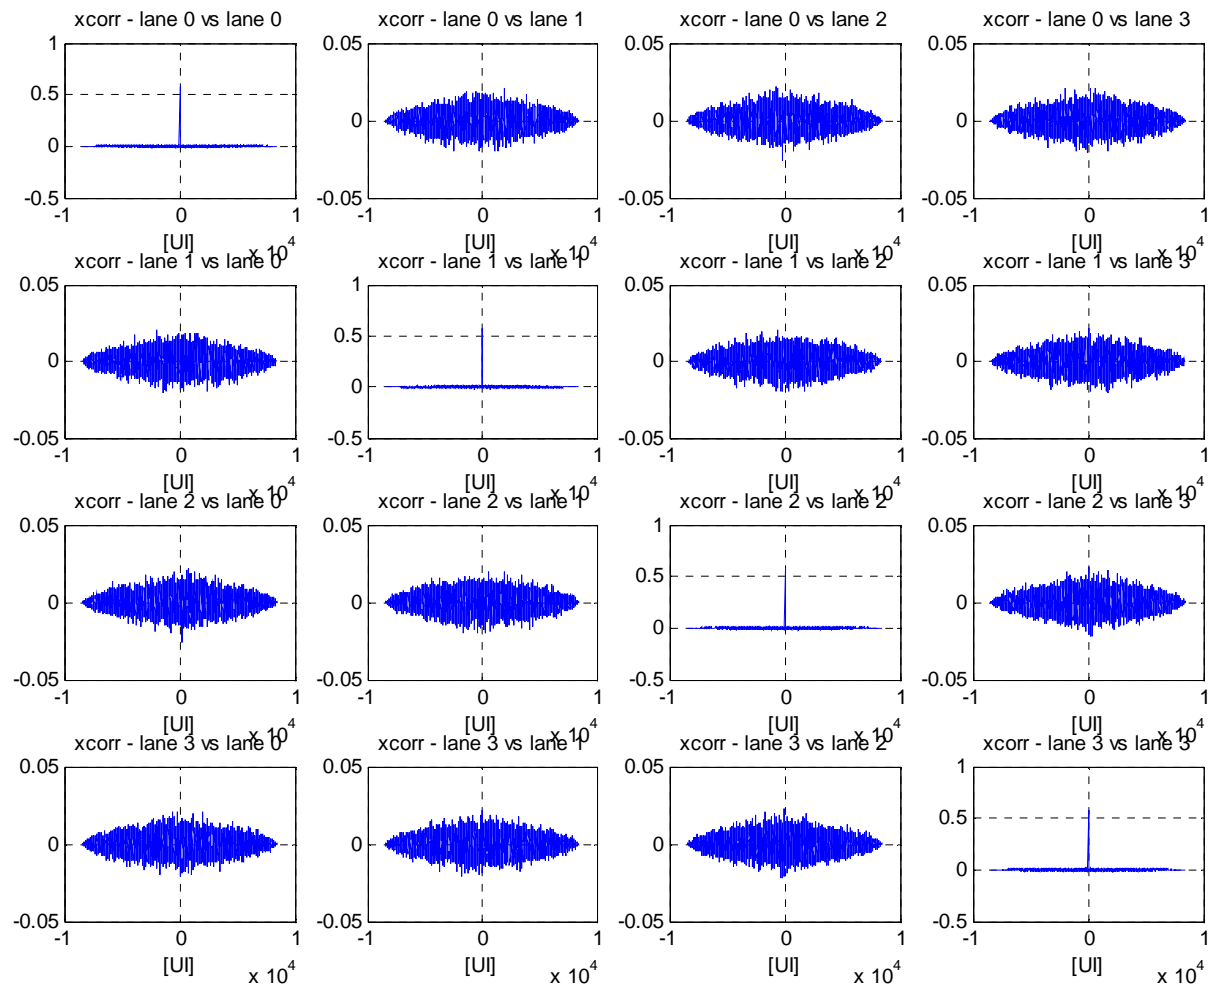
Block TX order

training frame word	45:26	35:26	25:16	15:6	5:0	
0	10'b1	10'b1	3'b1 7'b0	10'b0	6'b0	frame marker
DME	1111111111	1111111111	1110000000	0000000000	000000	
1	cell 15	cell 14	cell 13	cell 12	overhead	coef update
DME	1111111111	0000000000	1111100000	1111111111	000111	
value		0	0	1	0	1
2	cell 11	cell 10	cell 9	cell 8	overhead	coef update
DME	0000000000	1111111111	0000000000	1111111111	000111	
value		0	0	0	0	1
3	cell 7	cell 6	cell 5	cell 4	overhead	coef update
DME	0000000000	1111100000	1111111111	0000000000	111000	
value		0	1	0	0	1
4	cell 3	cell 2	cell 1	cell 0	overhead	coef update
DME	1111111111	0000000000	1111111111	0000000000	111000	
value		0	0	0	0	1

Good correlation properties

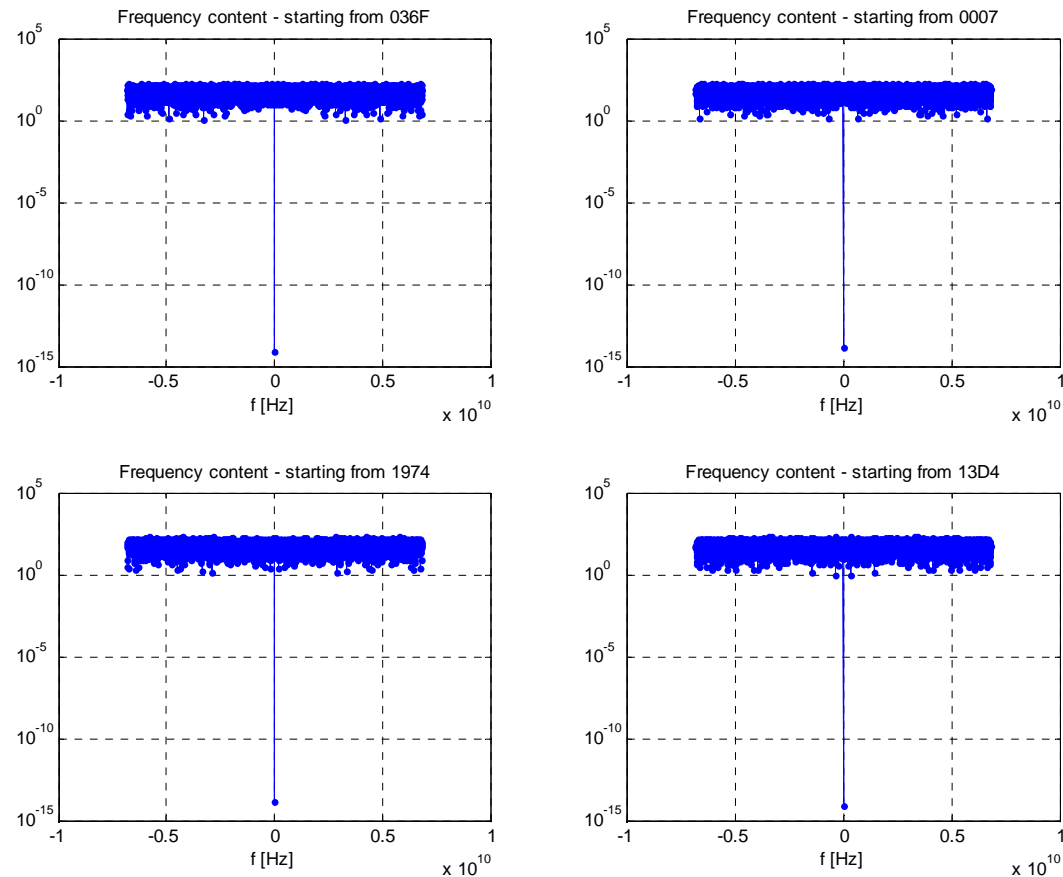
- After termination bits, gray coding, and pre-coding

these are not four rotations of the same sequence



Note: Training pattern without DME

Rich spectral content

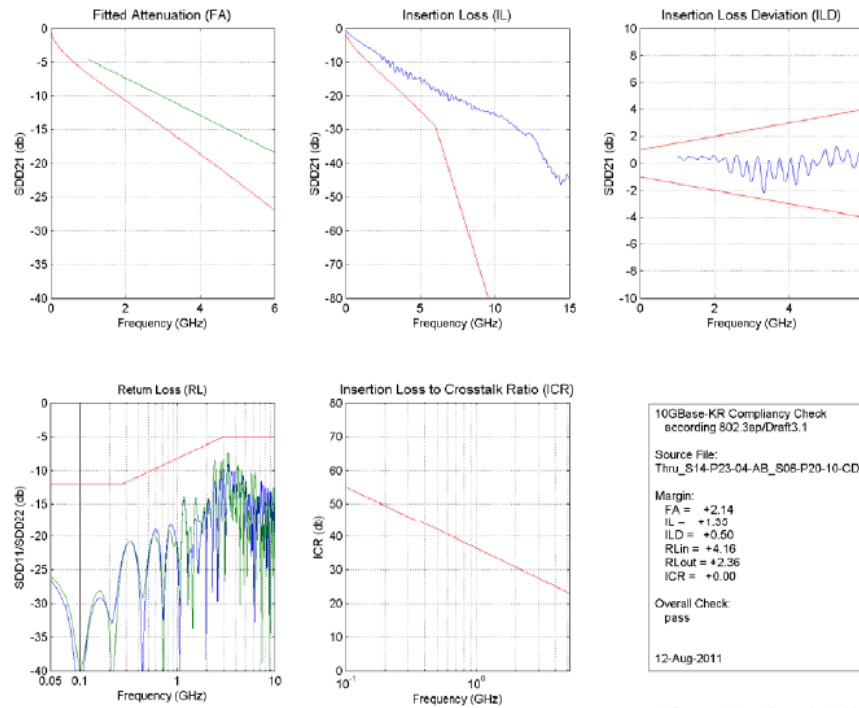


Note:
Training pattern
without DME

- The additional precoding and termination bit insertion create a distinct sequence for each lane.

Thru_S14-P23-04-AB_S06-P20-10-CD_NNN.s4p

Thru_S14-P23-04-AB_S06-P20-10-CD_NNN.s4p



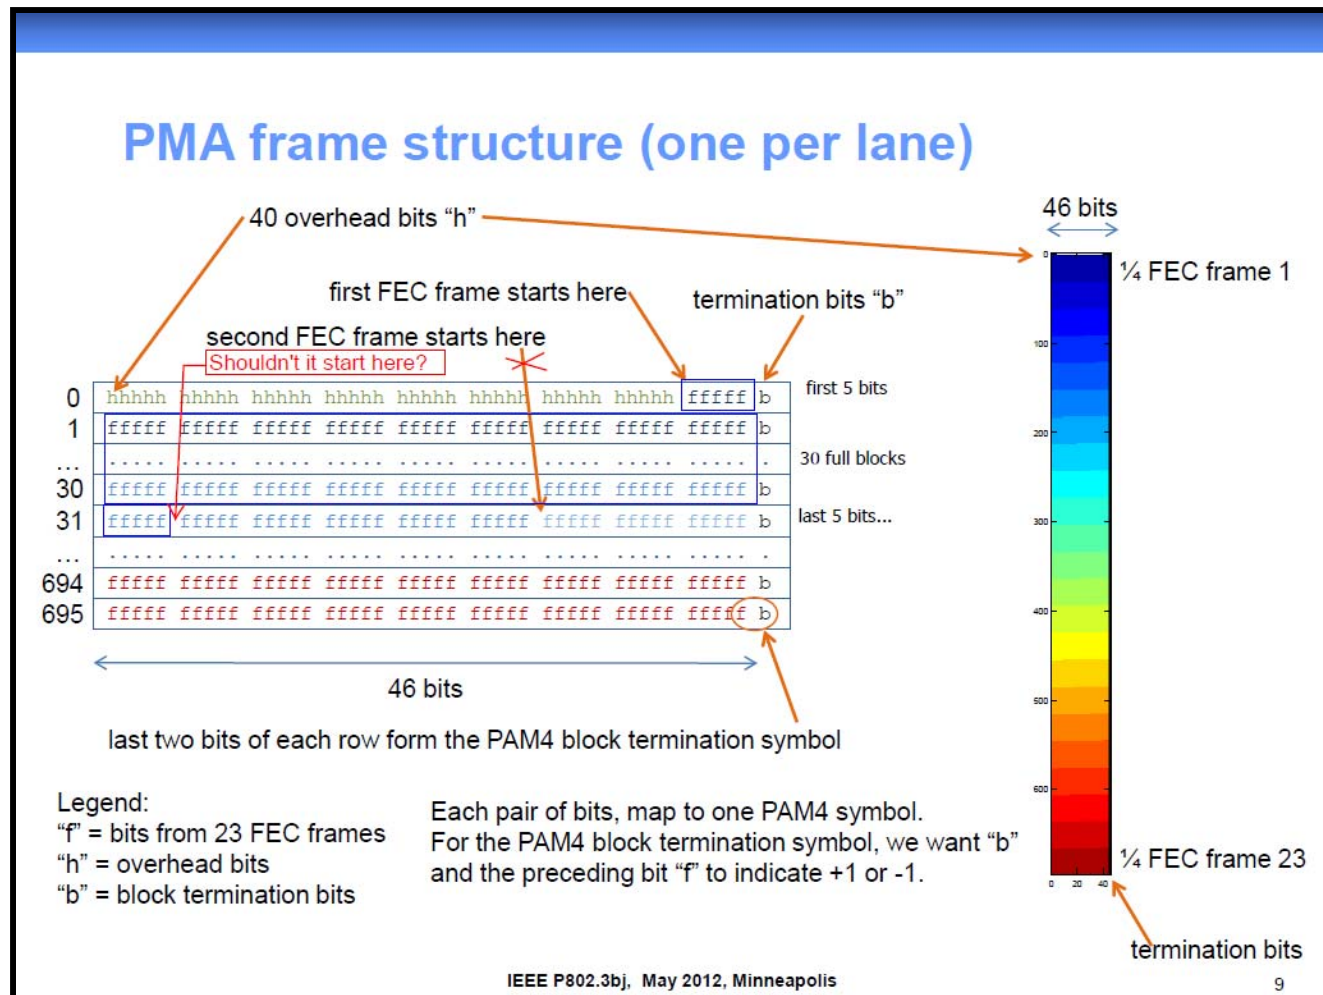
(c) Emerson Network Power GmbH 2011

Emerson Network Power

4



100GBASE-KP4 PMA Review



Taken from Brown_01a_0512.pdf (with corrections to 2nd frame starting location)

TFW to T-blocks to PAO Mapping

TF#	# T-blocks	PAO		TF#	# T-blocks	PAO
0	384	16		20	8064	17
1	768	3		21	8448	4
2	1152	19		22	8832	20
3	1536	6		23	9216	7
4	1920	22		24	9600	23
5	2304	9		25	9984	10
6	2688	25		26	10368	26
7	3072	12		27	10752	13
8	3456	28		28	11136	0
9	3840	15				
10	4224	2				
11	4608	18				
12	4992	5				
13	5376	21				
14	5760	8				
15	6144	24				
16	6528	11				
17	6912	27				
18	7296	14				
19	7680	2				

Training Frame Mapping

Block TX order



Symbol TX order

TFW	T-block	UI					
		45:26	35:26	25:16	15:6	5:0	
0	0:1	10'b1	10'b1	3'b1 7'b0	10'b0	6'b0	frame marker
1	2:3	cell 15	cell 14	cell 13	cell 12	overhead	coef update
2	4:5	cell 11	cell 10	cell 9	cell 8	overhead	coef update
3	6:7	cell 7	cell 6	cell 5	cell 4	overhead	coef update
4	8:9	cell 3	cell 2	cell 1	cell 0	overhead	coef update
5	10:11	cell 19	cell 18	cell 17	cell 16	overhead	status report
6	12:13	cell 15	cell 14	cell 13	cell 12	overhead	status report
7	14:15	cell 11	cell 10	cell 9	cell 8	overhead	status report
8	16:17	cell 7	cell 6	cell 5	cell 4	overhead	status report
9	18:19	cell 3	cell 2	cell 1	cell 0	overhead	status report
10	20:21	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
11	22:23	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
12-99	24:199	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
100	200:201	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
101	202:203	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
102	204:205	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
103-190	206:381	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern
191	382:383	prbs13	prbs13	prbs13	prbs13	prbs13	training pattern



PAM2

PAM4