

PAM4 transmitter training protocol

Adam Healey and Jeff Slavick

IEEE P802.3cd Task Force, July 2016

Contributors and supporters


Contributors to healey_3cd_01_0516

- Magesh Valliappan, Broadcom
- Vasu Parthasarathy, Broadcom

Supporters

- Eric Baden, Broadcom
- Matt Brown, Applied Micro
- Adrian Butter, Global Foundries
- Kent Lusted, Intel

Purpose of this presentation

- This is an update to [healey_3cd_01_0516](#) to address questions that were raised and provide additional detail in some areas
- Can precoding be used with PAM2 modulation? [Spoiler alert: No]
- How does precoding impact the training pattern properties?
- What are the settings for the transmitter equalizer at the beginning of training?
- Why include the receiver frame lock indication?
- Can we increase the time allowance for transmitter training (max_wait_timer)?
- New or updated slides are marked with a 

Introduction

- 10GBASE-KR (Clause 72) defines a start-up protocol that enables a receiver to control the equalizer of the remote transmitter
- This start-up protocol has been leveraged by subsequent PHYs operating over electrical backplanes to direct attach copper cable assemblies
- This is a useful feature that should be carried forward to PHYs based on 50 Gb/s per lane operating over backplane and copper cables
- This presentation proposes modifications to the protocol to make it more suitable for these next-generation applications

Assumptions and provisions

- PAM4 for 50 Gb/s per lane over backplanes and copper cables
- The transmitter equalizer will include additional taps
- Include a mechanism to select optional precoding (should it be added)




Transmitter training signal

Control and status (CS) fields

1 x 32 UI	Frame marker	16 x "3" followed by 16 x "0"
4 x 32 UI	Control	16 "cells", 8 UI per cell
	↓	Differential Manchester encoded (using "0" and "3" levels)
4 x 32 UI	Status	16 "cells", 8 UI per cell
	↓	Differential Manchester encoded (using "0" and "3" levels)

- Based on IEEE Std 802.3-2015, 72.10
- Support for PAM4 training pattern
- Higher ratio of training pattern to control/status



10GBASE-KR definition (for comparison):





Control and status fields

Control register encoding

Cell	Name	Description
15:14	Reserved	Transmit as 0, ignore on receipt
13:12	Initial condition request	<u>13 12</u> 1 1 = Preset 3 (TBD) 1 0 = Preset 2 (TBD) 0 1 = Preset 1 (no equalization) 0 0 = Individual coefficient control
11:10	Reserved	Transmit as 0, ignore on receipt
9:8	Modulation and precoding request*	<u>9 8</u> 1 1 = PAM4 with precoding 1 0 = PAM4 0 1 = Reserved 0 0 = PAM2/NRZ
7:5	Reserved	Transmit as 0, ignore on receipt
4:2	Coefficient select	<u>4 3 2</u> 1 1 0 = $c(-2)$ 1 1 1 = $c(-1)$ 0 0 0 = $c(0)$ 0 0 1 = $c(1)$
1:0	Coefficient request	<u>1 0</u> 1 1 = No equalization (0) 1 0 = Decrement 0 1 = Increment 0 0 = Hold

Status register encoding


Cell	Name	Description
15	Receiver ready	1 = Training is complete and the receiver is ready for data 0 = Request for training to continue
14:12	Reserved	Transmit as 0, ignore on receipt
11:10	Modulation and precoding status*	<u>11 10</u> 1 1 = PAM4 with precoding 1 0 = PAM4 0 1 = Reserved 0 0 = PAM2/NRZ (default)
9	Receiver frame lock	1 = Frame boundaries identified 0 = Frame boundaries not identified
8	Initial condition status	1 = Updated 0 = Not updated
7:5	Reserved	Transmit as 0, ignore on receipt
4:2	Coefficient select echo	<u>4 3 2</u> 1 1 0 = $c(-2)$ 1 1 1 = $c(-1)$ 0 0 0 = $c(0)$ 0 0 1 = $c(1)$
1:0	Coefficient status	<u>1 0</u> 1 1 = Coefficient not supported 1 0 = Coefficient at limit 0 1 = Updated 0 0 = Not updated

* Modulation and precoding apply to the training pattern only.

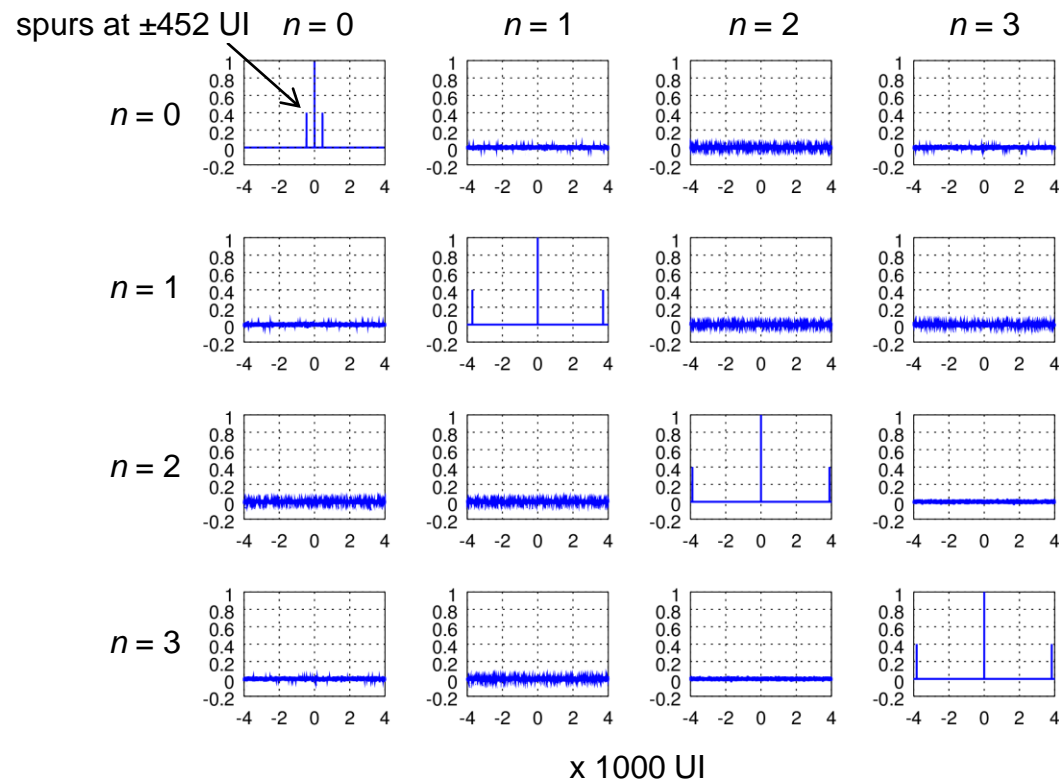
Training pattern

n	Polynomial $_n, G(x)$	Notes
0	$1 + x + x^2 + x^{12} + x^{13}$	PRBS13Q polynomial
1	$1 + x^2 + x^3 + x^7 + x^{13}$	
2	$1 + x^2 + x^4 + x^8 + x^{13}$	
3	$1 + x^2 + x^5 + x^9 + x^{13}$	

Generator example ($n = 0$)



NOTE: A is the bit arriving first

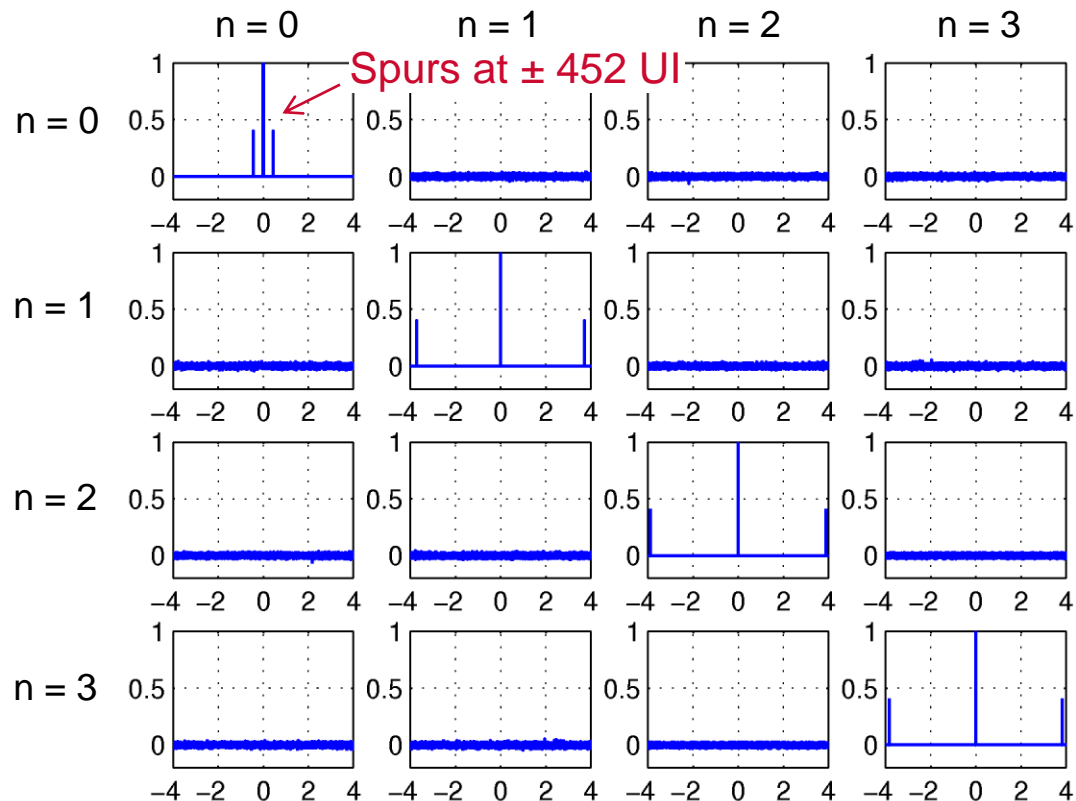


- Based on IEEE Std 802.3-2015, 92.7.12
- Configurable generator polynomial and seed to avoid correlated interference between lanes
- When modulation status is PAM2, map A (or B) bits to “0” and “3” levels
 - A and B are shifted copies of the PRBS13 sequence
- Gray-mapped PAM4 symbols may optionally be precoded

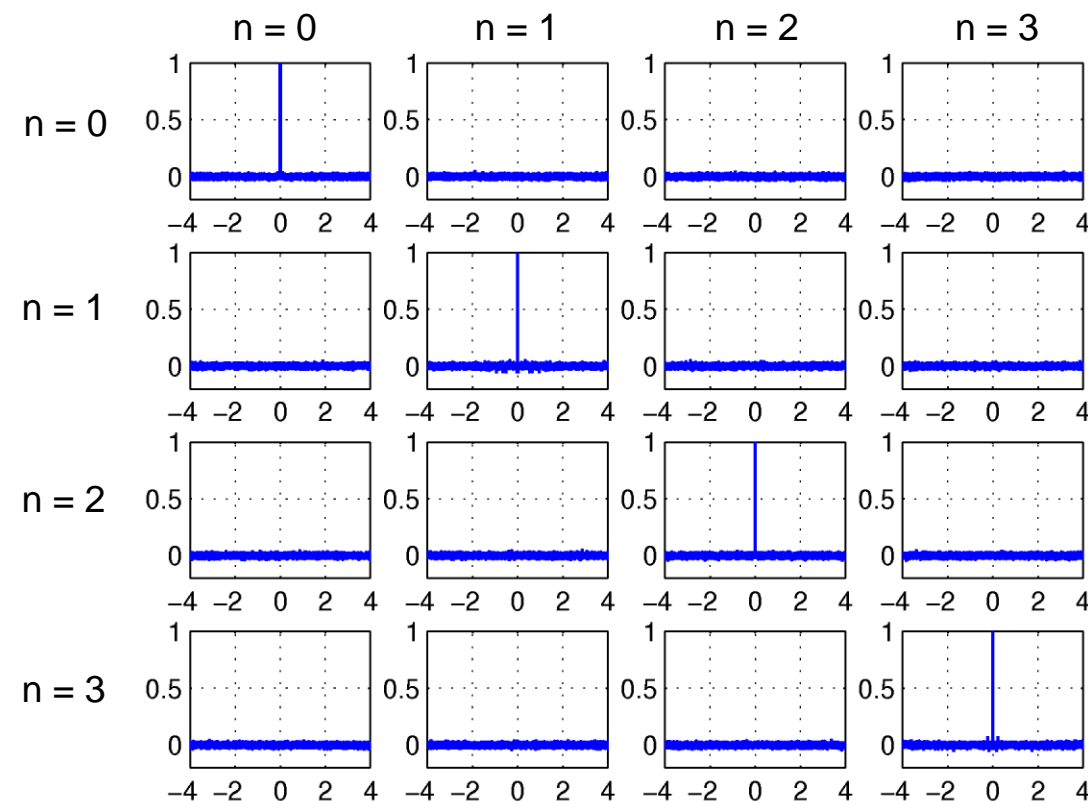


Cross-correlations without and with precoding

PAM4



PAM4 with precoding





Training pattern properties without and with precoding

Property	PAM4				PAM4 with precoding			
	0	1	2	3	0	1	2	3
Symbol count								
0	2047	2047	2047	2047	2027	2057	2035	2119
1	2048	2048	2048	2048	2081	2021	2050	2044
2	2048	2048	2048	2048	2069	2039	2061	1977
3	2048	2048	2048	2048	2104	2074	2045	2051
Baseline wander, pk-pk [1]	2.6	1.9	2.4	2.3	2.4	1.9	2.2	3.4
Transition density, min. [2]								
All transitions	73.5	73.4	73	72.9	72.7	73.6	73.2	73
Transitions through 0 V	48.2	48	48.1	47.5	46.3	47.6	48.5	48.1
All symmetric transitions	48.2	48	48.1	47.5	48.2	48	48.1	47.5
All symmetric transitions through 0 V	23	23.3	23.7	23.4	23	23.3	23.7	23.4
Longest fully-represented sequence [3]	6	6	6	6	5	5	5	5

[1] Normalized to twice the transmitter steady-state output amplitude. Measured with high-pass corner frequency set to 1/10,000 of the signaling rate.

[2] Measured with the corner frequency set to 1/6,640 of the signaling rate.

[3] When a sequence of length n is “fully-represented”, all L^n combinations of L-level symbols appear in the test pattern.

Principles of operation

- Define only the response to defined requests
- Use of defined requests is beyond the scope of the standard
- Frame lock diagram similar to Figure 72-4 accounting for different frame marker intervals
- Training state diagram similar to Figure 72-5
- New coefficient update state diagram to account for initial condition control and individual coefficient processing

Coefficient update state diagram

coef_req: Enumerated variable derived from the coefficient request bits from control field of the received training frames. This variable may be one of the following values: hold, decrement, increment, no equalization.

coef_sel: Variable derived from the coefficient select bits from the control field of the received training frames. It is assigned a signed integer value that is the 2's complement interpretation of the bits.

coef_sts: Enumerated variable that may be assigned one of the following values: not updated (not_upd), updated, coefficient at limit, coefficient not supported.

ic_req: Enumerated variable derived from the "initial condition request" bits from the control field of the received training frames. This variable may be one of the following values: individual control (ind_ctl), preset 1, preset 2, preset 3.

ic_sts: Enumerated variable that may be assigned one of the following values: not updated (not_upd), updated.

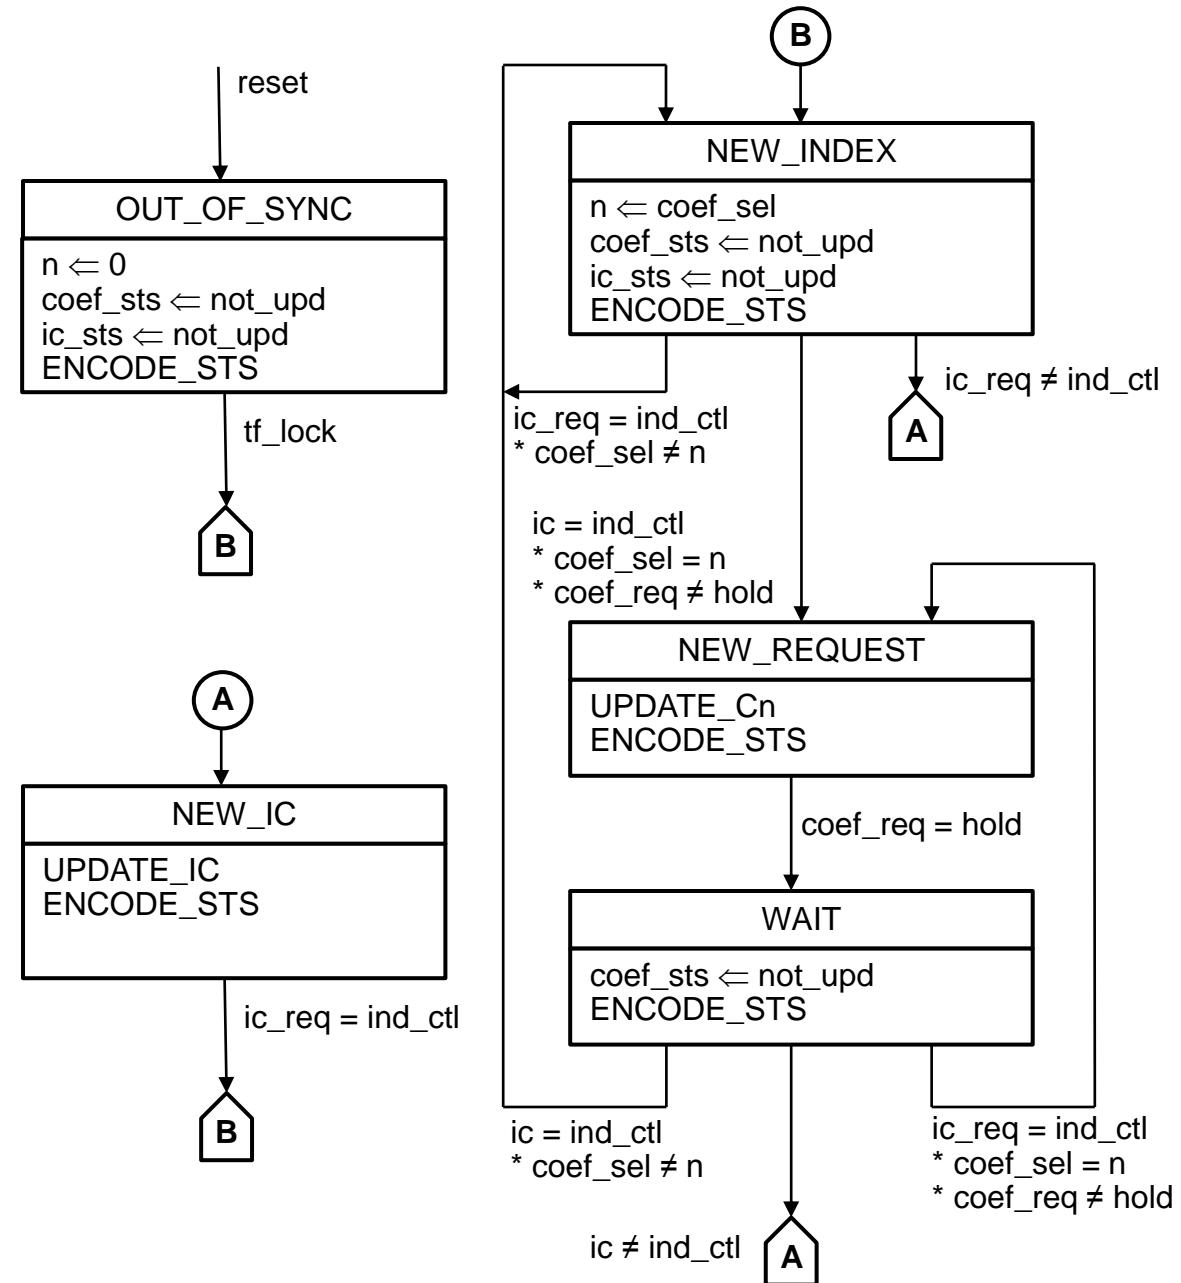
n: Variable that stores the most recent value of coef_sel.

tf_lock: Boolean variable that is true when the training frame marker positions have been identified and is false otherwise.

ENCODE_STS: Function that encodes portions of the status field of transmitted training frames. The variable tf_lock is mapped to the receiver frame lock bit, n is mapped to the coefficient select echo bits, coef_sts is mapped to the coefficient status bits, and ic_sts is mapped to the initial condition status bit.

UPDATE_Cn: Function that updates the value of $c(n)$ based on the current values of n and coef_req. The result of the update is stored in the variable coef_sts.

UPDATE_IC: Function that updates all equalizer coefficients based on the current value of ic_req. The result of the update is stored in the variable ic_sts.



Coefficient update state diagram, continued

UPDATE_Cn algorithm:

```
if n in n_list

    if coef_req = increment
         $c(n) = c(n) + cn\_stp$ 
    else if coef_req = decrement
         $c(n) = c(n) - cn\_stp$ 
    else if coef_req = no equalization
         $c(n) = 0$ 
    else
         $c(n) = c(n)$ 

    if  $c(n) > cn\_max$ 
         $c(n) = cn\_max$ 
        coef_sts = coefficient at limit
    else if  $cn < cn\_min$ 
         $c(n) = cn\_min$ 
        coef_sts = coefficient at limit
    else
        coef_sts = updated

else
    coef_sts = coefficient not supported
```

UPDATE_IC algorithm:

```
if ic_req = preset 1
    <set coefficients to preset 1>
else if ic_req = preset 2
    <set coefficients to preset 2>
else if ic_req = preset 3
    <set coefficients to preset 3>
else
    <no action>

ic_sts = updated
```

n_list: Set of valid coefficient indices
c(n): The current value of the coefficient n
cn_stp: Step size for coefficient c(n)
cn_min: Minimum value for coefficient c(n)
cn_max: Maximum value for coefficient c(n)

- Functions are called once upon entry into a state sequentially in the order listed



Additional principles of operation

- ~~The default modulation for the training pattern is PAM2/NRZ~~
- ~~Precoding is off by default~~
- ~~Precoding mode is independent of the mode chosen for the opposite direction of the link~~
- ~~Receiver may request changes to modulation (or precoding) at any point during the training process~~
- ~~Modulation (and precoding) apply to the training pattern only; encoding of frame marker, control field, status field, and “00” pad are unchanged~~
- ~~The precoding mode at the end of transmitter training is used for “mission” data~~
- ~~When the “Receiver frame lock” flag is 1, the receiver shall respond to requests within TBD ms (2 ms in 92.7.12)~~



Additional principles of operation

- The default modulation for the training pattern is PAM2/NRZ
- The default transmitter equalization is “Preset 1 (no equalization)”
- Modulation/precoding mode is independent of the mode chosen for the opposite direction of the link
- Receiver may request changes to modulation/precoding at any point during the training process
- Training must end with modulation/precoding set to either “PAM4” or “PAM4 with precoding” otherwise failure will be indicated
- If training completes with “modulation and precoding status” set to “PAM4 with precoding”, then the subsequent “mission” data will also be precoded



Additional principles of operation, continued

- Changes to modulation/precoding may occur at any time between the point the request is received and the point the change is acknowledged
- Modulation/precoding control applies to the training pattern only; encoding of the frame marker, control field, status field, and “00” pad are unchanged
- The precoder operates on the first PAM4 symbol of the training pattern as if the previous output symbol was “0”
- When the “Receiver frame lock” flag is 1, the receiver shall respond to requests within TBD ms (2 ms in 92.7.12)



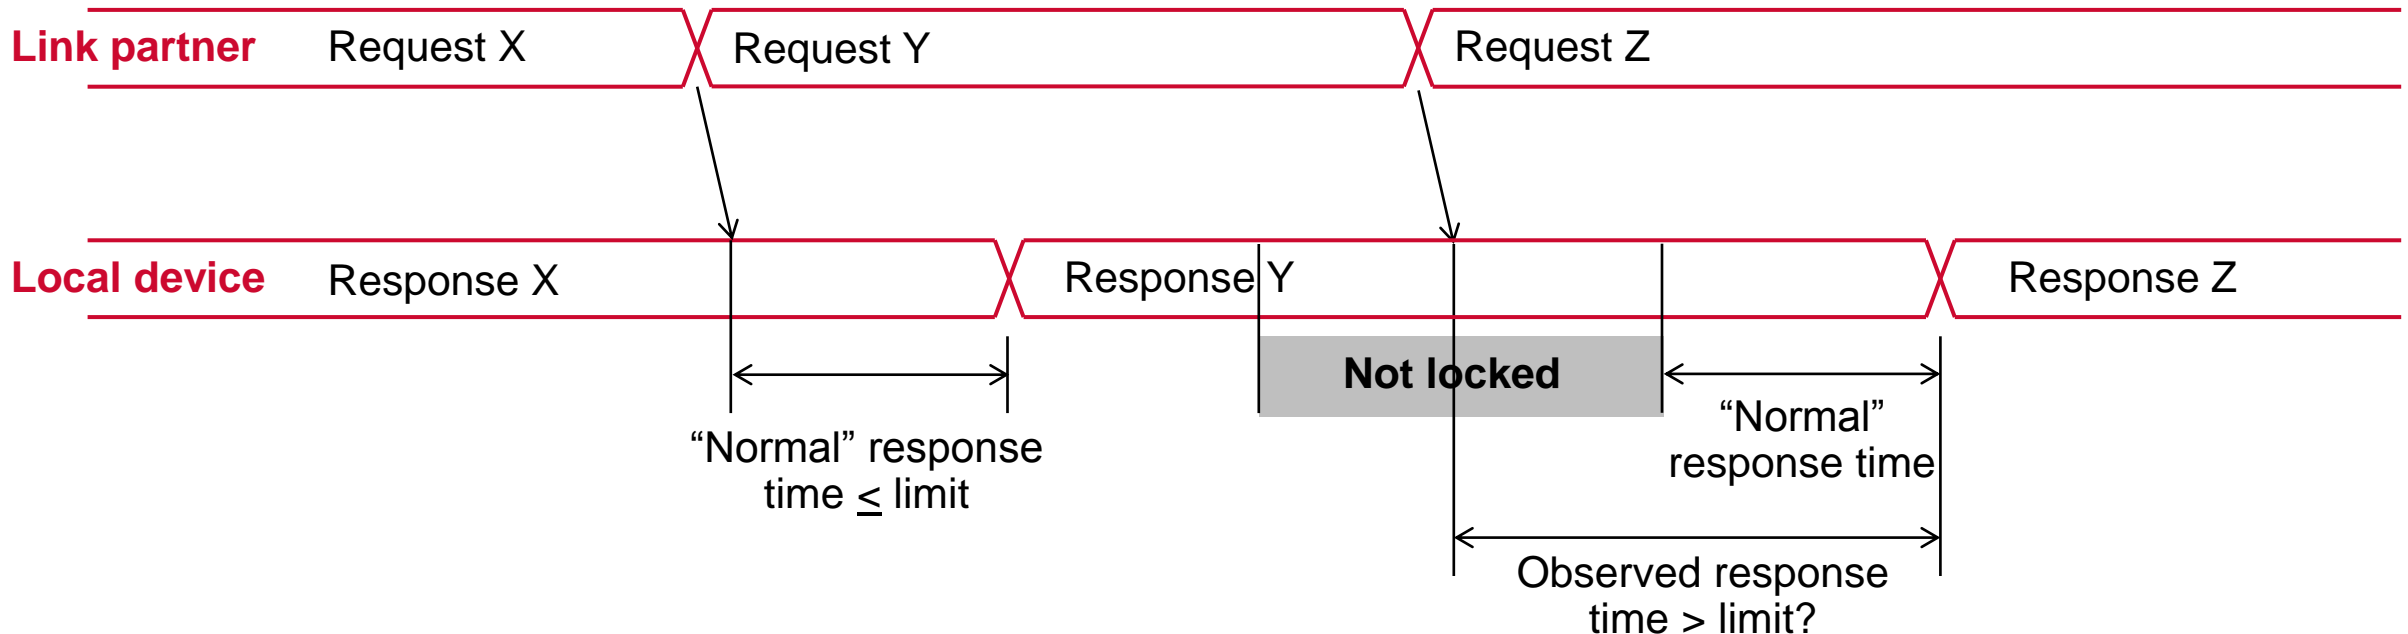
Why include the receiver frame lock indication?

- IEEE Std 802.3-2015, 92.7.12 item b)
 - *In addition to the coefficient update process specified in 72.6.10.2.5, the period from receiving a new request to responding to that request shall be less than 2 ms, except during the first 50 ms following the beginning of the start-up protocol. The beginning of the start-up protocol is defined to be entry into the AN_GOOD_CHECK state in Figure 73–10.*
- “The first 50 ms” is the allowance for the receiver to lock to the training frame
- After this grace period, the 2 ms response time applies to the receiver whether or not it is actually locked
- Failure to lock within the allowed time, or [temporarily] losing lock after the initial allowance passed, could result in a violation of the response time requirement
- The violation is ambiguous; is the receiver too slow to respond or did it simply not get the message in time?



Example

- IEEE Std 802.3-2015, 92.7.12 item b)
 - The start of the period is the frame marker of the training frame with the new request and the end of the period is the frame marker of the training frame with the corresponding response.
 - A new request occurs when the coefficient update field is different from the coefficient field in the preceding frame. The response occurs when the coefficient status report field is updated to indicate that the corresponding action is complete.





Determining whether or not the receiver is locked

- The value of the frame_lock variable is presented in the BASE-R PMD register (e.g., 1.151.1, see 45.2.1.81.2)
- However, per IEEE Std 802.3-2015, 45.1
 - *The MDIO electrical interface is optional. Where no physical embodiment of the MDIO exists, provision of an equivalent mechanism to access the registers is recommended.*
- It is possible to determine if the receiver is locked using MDIO (or an equivalent interface), but such an interface is not always available
- If the frame_lock value is included in the status register, it may be observed even when MDIO is not available
- This enables a more straightforward and verifiable response time requirement



Comparison of the two schemes

92.7.12 item b)

- To be compliant, an implementation must achieve lock within 50 ms regardless of operating conditions
- An implementation must plan for the possibility that the first 50 ms of training will be spent waiting for the link partner to achieve lock
- If frame lock is lost after the first 50 ms of training, the receiver must be able to reacquire lock and respond to a new request within the allocated response time

Proposed scheme

- No arbitrary requirement to achieve or maintain lock
- Local device is aware of when the link partner acquires frame lock and can begin training at that time (assuming it is also locked)
- Response time requirements recognize the current state of the receiver



Can we increase the time allowance for transmitter training?

- The Clause 72 protocol allocates approximately 500 ms for transmitter training
- There is interest in increasing this time for 50 Gb/s backplane and copper cable links using PAM4
- Propose to increase the time to approximately 1.5 seconds for PAM4 links
- The additional time could be well used and still results in reasonable “link up” times
- This requires an increase in the max_wait_timer (see 72.6.10.3.2) specifically for this protocol
- This also requires an increase to the Auto-Negotiation link_fail_inhibit_timer for PHYs that use this protocol
- This may also enable a longer (than 2 ms) response time requirement if desired

Management provisions

- Local device (LD) and link partner (LP) control and status registers (per lane)
- PMD training pattern register: polynomial identifier, seed (per lane)
- Enable and restart training control bits (per lane)
- Frame lock, start-up protocol status, training failure status bits (per lane)

Summary and conclusions

- The Clause 72 start-up protocol is a useful feature that should be carried forward to PHYs based on 50 Gb/s per lane operating over backplane and copper cables
- This presentation proposes modifications to the protocol to make it more suitable for these next-generation applications
 - Support for **PAM4** training patterns
 - Higher ratio of training pattern to control/status
 - Support for **additional transmitter equalizer taps**
 - Enable selection of optional **precoding**