

Proposal for 10Gb/s single-lane PHY using PAM-4 signaling

Rob Brink, Agere Systems
Bill Hoppin, Synopsys

Supporters

Ted Rado, Analogix

John D'Ambrosia, Tyco Electronics*

* This contributor supports multi-level signaling standardization for certain applications. This support does not necessarily reflect the support of PAM-4 over competing technology solutions.

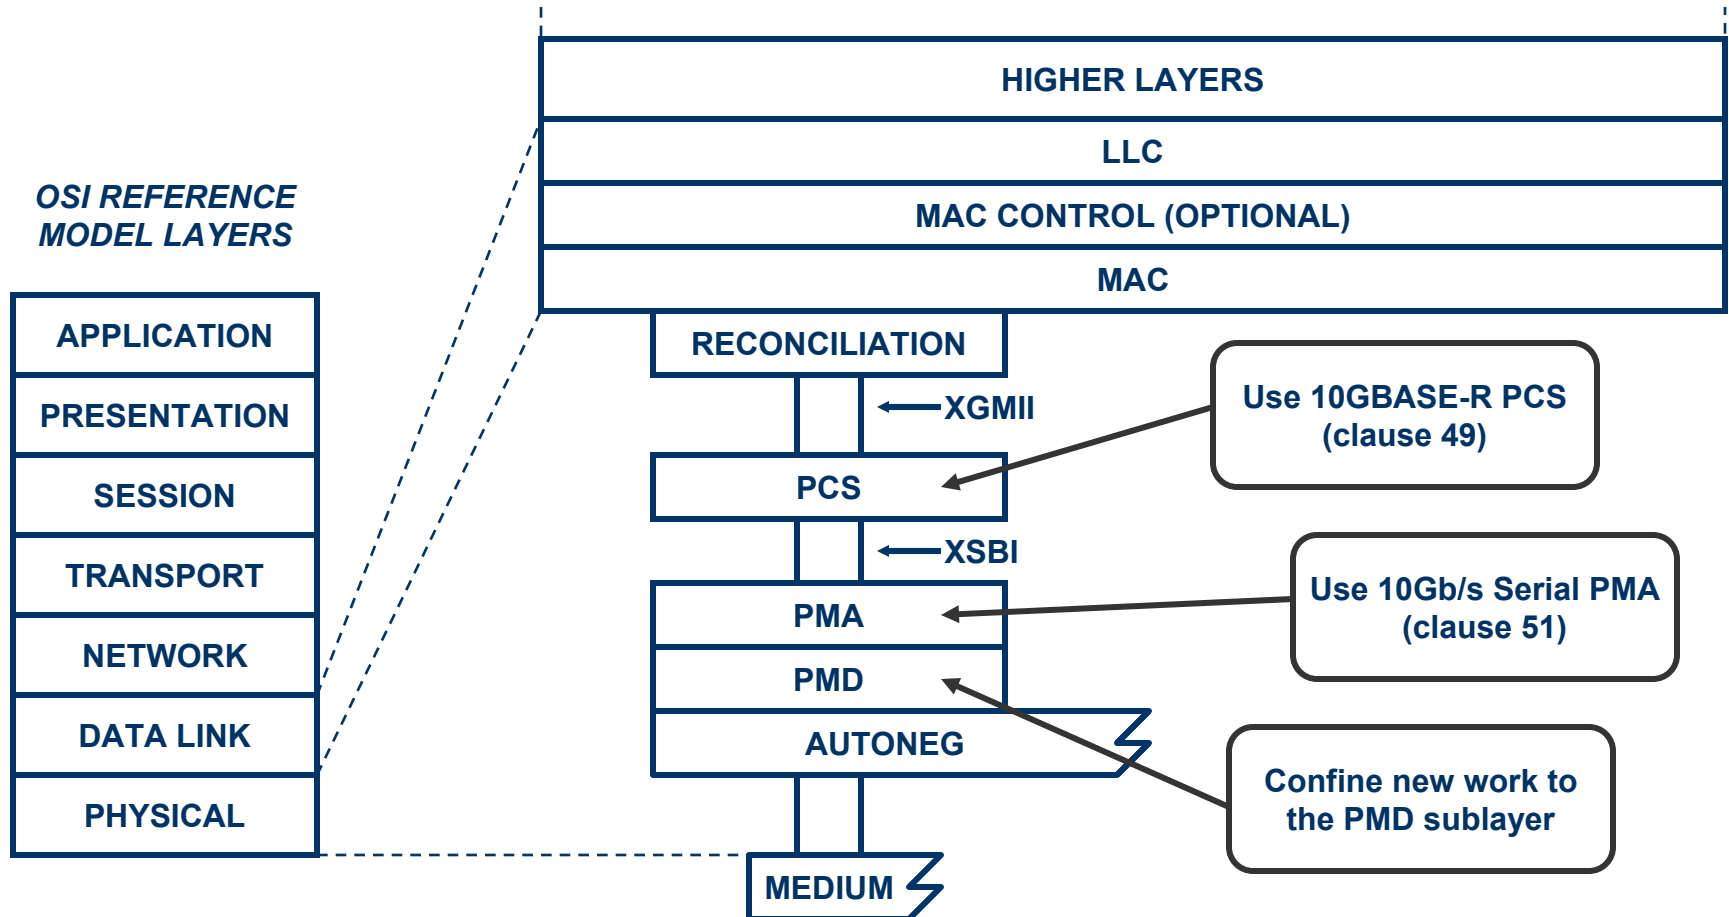
Scope and Purpose

- This presentation proposes a new PMD sublayer based on PAM-4 signaling.
- The new PMD leverages the 10GBASE-R PCS (clause 49) and 10Gb/s serial PMA (clause 51) to form a complete physical layer stack.
- This presentation describes the fundamental concepts behind the proposed PMD.
- This presentation describes how the proposed PMD satisfies the Task Force objectives for the single-lane 10Gb/s PHY.

Agenda

- **Proposal Overview**
- Link Simulations
- Link Initialization Protocol (LIP) Detail
- Conclusions

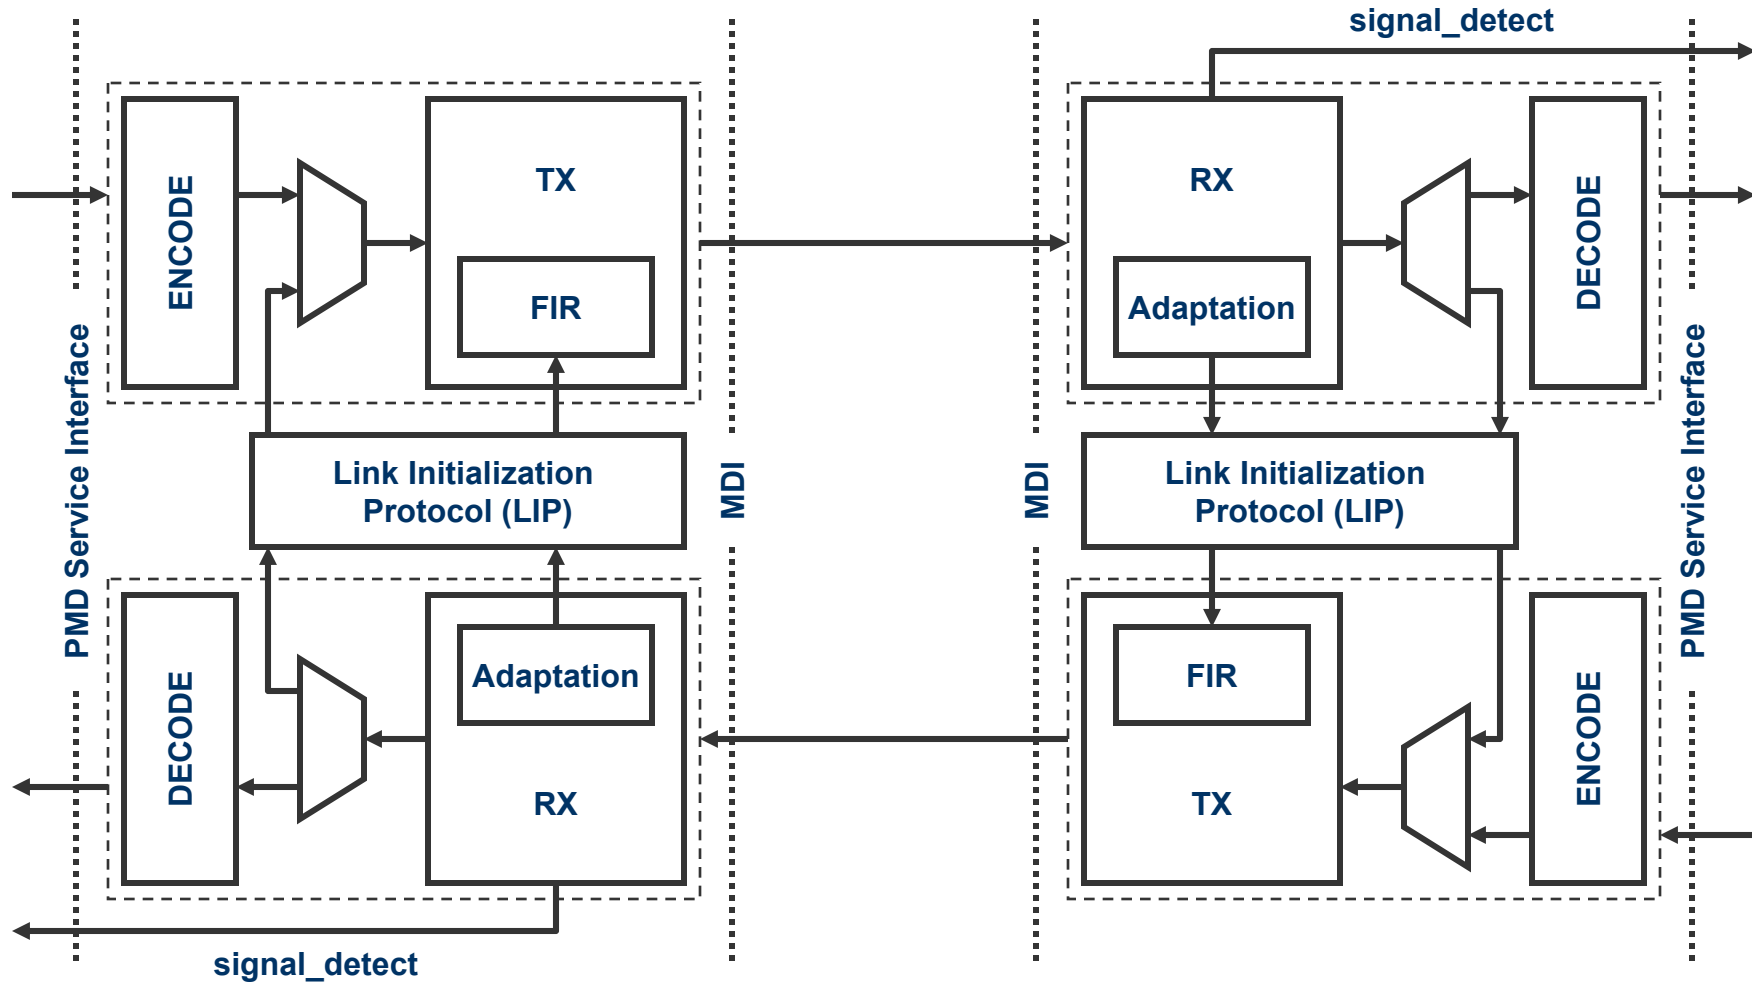
Layer Model



Proposal Overview

- Reduce occupied bandwidth through the use of PAM-4 signaling.
 - Reduces required equalization effort.
 - SNR improvement for worst-case channel exceeds the 9.5dB lost to multi-level.
- Divide equalization effort between the transmitter and receiver.
- Define an adaptive transmitter.
 - Precise equalization is easier to implement at the transmitter.
 - Alleviates burden on receiver circuitry.
 - Transmitter is trained during link initialization, and then the settings are frozen.
 - Requires a receiver-to-transmitter communication path (but only during link initialization).
- Continuously adaptive receiver.
 - Simpler, lower power design due to pre-compensation at the transmitter.
 - Tracks time-variation due to temperature and humidity changes.

Link Model



Encoding/Decoding

- Each PAM-4 symbol carries two information bits.
 - 10.3125Gb/s → 5.15625Gbaud
- Simple linear encoding preserves DC balance.

Bits [MSB, LSB]	Symbol
00	-3
10	-1
01	+1
11	+3




Agenda

- Proposal Overview
- **Link Simulations**
- Link Initialization Protocol (LIP) Detail
- Conclusions

Link Simulations

- Basic premise and feasibility are demonstrated using channel data representative of the worst-case environment.
- Transmitter contains 5-tap adaptive finite impulse response (FIR) filter.
 - Filter is trained using only -3 and +3 symbols, as described later.
 - Training pattern is PN-7.
 - For the purpose of this simulation, LMS adaptation is employed.
- Receiver equalizer is modeled as a simple gain peaking amplifier.
 - No time varying element in this simulation.
- Following training, the PAM-4 eye is evaluated.
- Sample point is positioned at eye center.
- Vertical and horizontal eye opening is reported at 1E-15.


Test Channel




Note: Tx/Rx load model not intended to represent a specific implementation. Rather, its purpose is to ensure that mismatch effects are included in the simulation.

Equalizer Training




Transmitter
Gaussian Pulse, $T_r(20-80\%) = 50\text{ps}$





Channel



Receiver
4dB Gain Peaking at 2.5GHz





10Gb/s Operation ($0.05U_{I_{p-p}}$ Tx RJ, no crosstalk)





Vertical Eye Opening (at 1E-15)	0.16	au
Horizontal Eye Opening (at 1E-15)	0.28	UI
Effective DJ, Peak-Peak	0.65	UI
Effective RJ, RMS	0.004	UI

Crosstalk

- Only single-aggressor NEXT and single-aggressor FEXT applied.
 - Exceeds proposed multi-disturber allocation.
- Near-end aggressors assumed to be asynchronous with respect to the signal of interest ($\pm 100\text{ppm}$).
 - Peak value “walks” across eye.
- Far-end aggressors assumed to be synchronous with respect to the signal of interest.
 - Peak value fixed at eye center (worst-case analysis).
- Near-end and far-end aggressors assumed to be similar transmitters driving similar channels.
 - Same output amplitude, rise time, and FIR tap settings.





10Gb/s Operation ($0.05U_{I_{p-p}}$ Tx RJ, crosstalk)



Vertical Eye Opening (at 1E-15)	0.10	au
Horizontal Eye Opening (at 1E-15)	0.19	UI
Effective DJ, Peak-Peak	0.66	UI
Effective RJ, RMS	0.010	UI

Jitter

- In this simulation, deterministic jitter is the intrinsic jitter due to unconstrained switching among PAM-4 levels.
- Random jitter is increased from base value 0.05 to $0.15U_{p-p}$ (as measured at $1E-15$).
- Note that at $0.10U_{p-p}$, a $1000mV_{ppd}$ output voltage will yield at $45mV_{ppd}$ eye opening at the slicer input.




NOTE: In this simulation, eye height is normalized to a $2V_{p-p}$ Tx output voltage. This does not imply that the solution requires $2V_{p-p}$.

Aside: NRZ Link Simulations


- Driver rise time and Tx / Rx termination models changed to be more appropriate for a 10Gb/s NRZ design.
- Transmitter contains 3-tap adaptive finite impulse response (FIR) filter (two pre-cursor taps).
- Receiver is modeled as a gain peaking filter followed by a 5-tap decision feedback equalizer.
 - Gain peaking at $f_{\text{baud}}/2$ is identical to PAM-4 gain peaking at $f_{\text{baud}}/2$.
- Transmit and receive equalizers are jointly trained.
 - Training pattern is PN-7, LMS adaptation is employed.
- Following training, the NRZ eye is evaluated.
- Sample point is positioned at eye center
 - ...as seen at the output of the gain peaking filter.
- Vertical and horizontal eye opening is reported at 1E-15.

Equalizer Training (NRZ)





Transmitter
Gaussian Pulse, $T_r(20-80\%) = 30\text{ps}$





Channel



Receiver
4dB Gain Peaking at 5GHz




NRZ Eye (0.05UI_{p-p} Tx RJ, no crosstalk)



Vertical Eye Opening (at 1E-15)	0.16	au
Horizontal Eye Opening (at 1E-15)	0.45	UI
Effective DJ, Peak-Peak	0.48	UI
Effective RJ, RMS	0.004	UI

Identical to PAM-4 eye height!

Note: Crosstalk completely closes the eye at 1E-15



Agenda

- Proposal Overview
- Link Simulations
- **Link Initialization Protocol (LIP) Detail**
- Conclusions

Link Initialization Protocol (LIP)

- Facilitates clock recovery.
- Optimizes transmitter FIR.
- Automatic power control.
 - Receiver may steer the transmitter output voltage to the minimum level required for acceptable performance.
- Optimize receiver equalizer.

LIP Frame Format

- Transmitted using only -3 and +3 symbols.
 - NRZ signaling at 5.15625Gb/s.
- Frame length is 560 bits.
 - Divisible by both 16 and 20.
 - 4-byte frame marker, 8-byte control channel, 58-byte training pattern

Transmission Order



Frame Marker


- Delimits LIP frames.
- Fixed 4-byte pattern, 0xFFFF_0000
 - Detectable over unequalized or partially equalized channels.
 - Does not occur in control channel or training pattern.
 - Also may be used as a polarity check (reception of 0x0000_FFFF indicates polarity reversal).

Control Channel

- 2-bytes of control information (8-bytes after encoding).
 - Status report.
 - Coefficient update.
- Double-Wide Manchester Coding
 - Guarantees 50% transition density.
 - Guarantees DC balance.
 - Prevents frame marker pattern from appearing in the control channel.
 - Detectable over unequalized or partially equalized channels.

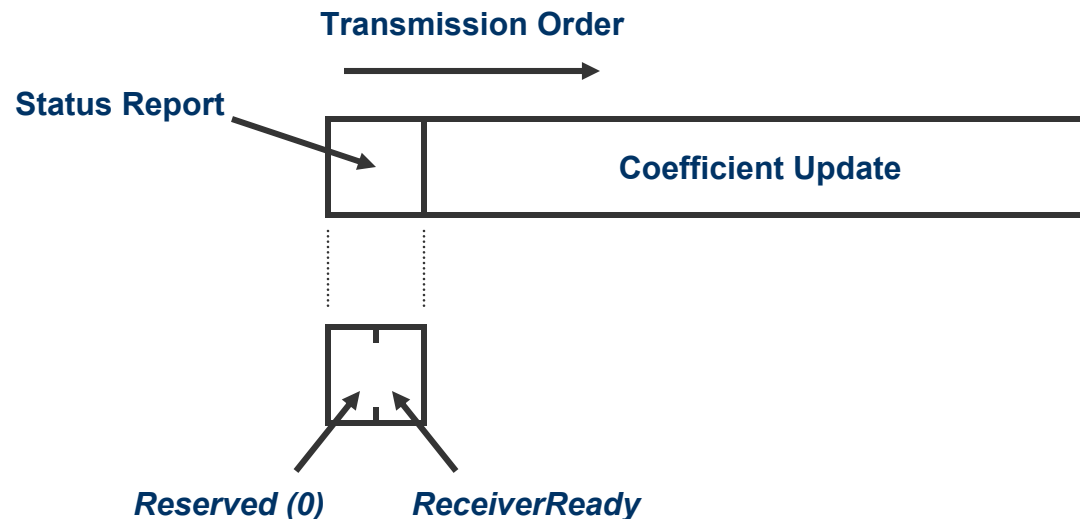
Message Bit	Encoded Sequence
0	1100
1	0011

Transmission Order



Status Report

- *ReceiverReady* indicator (1-bit).
 - Asserted (1) when receiver deems that equalization training (for both the transmitter and receiver) is complete.



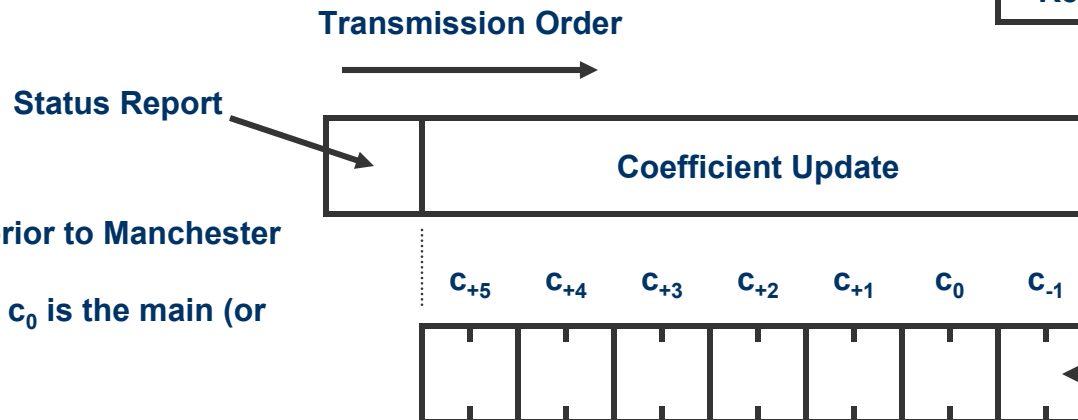
Notes

- a) Fields shown prior to Manchester encoding.

Coefficient Update

- Supports parallel update of transmitter FIR coefficients to a maximum of 7 taps.
 - It is not necessary for an implementation to support all 7 taps.
- Each tap has an associated action.
 - Decrement / Hold / Increment
 - Agnostic to the supported tap weight resolution.
 - Tolerant of corrupted or lost coefficient updates.
 - Actions applied to unsupported taps are ignored.

Action	Encoding
Hold	00
Decrement	01
Increment	10
Reserved	11



Notes

- Fields shown prior to Manchester encoding.
- By convention, c_0 is the main (or gain) tap.

Training Pattern

- Any DC-balanced “random” pattern will suffice.
- One possibility is the 464-bit pattern consisting of the pattern shown below (232-bits) followed by its inverse.

Transmission Order
→

Sync. Pattern [6-bytes]	00 FF 00 FF 00 FF
Impulse [3-bytes]	00 80 00
High-Speed Clock [4-bytes]	AA AA AA AA
1'b1 followed by $x^7 + x^6 + 1$ (all ones seed) [16-bytes]	FE 04 18 51 E4 59 D4 FA 1C 49 B5 BD 8D 2E E6 55

↓

LIP Highlights (1/2)

- LIP frames are signaled continuously using only -3 and +3 symbols.
 - Absence of -1 and +1 symbols for an extended period indicates that the remote PMD wishes to re-initialize.
- Local receiver adaptation process sends FIR tap weight updates to the remote transmitter via the coefficient update field.
 - The adaptation process itself is beyond the scope of the standard.
 - A variety of algorithms may be employed.

LIP Highlights (2/2)

- When the local adaptation process determines that the local Tx and remote Rx are fully trained, it sets the ReceiverReady bit on outgoing LIP frames.
 - The LIP state machine must see the ReceiverReady bit asserted three consecutive times before it concludes that remote receiver is ready to receive data (no hair triggers).
- When the LIP state machine determines that the local and remote receivers are ready to receive data, it sends a fixed number of LIP frames to ensure that the remote receiver properly detects the ReceiverReady bit.

LIP State Diagram (1/3)

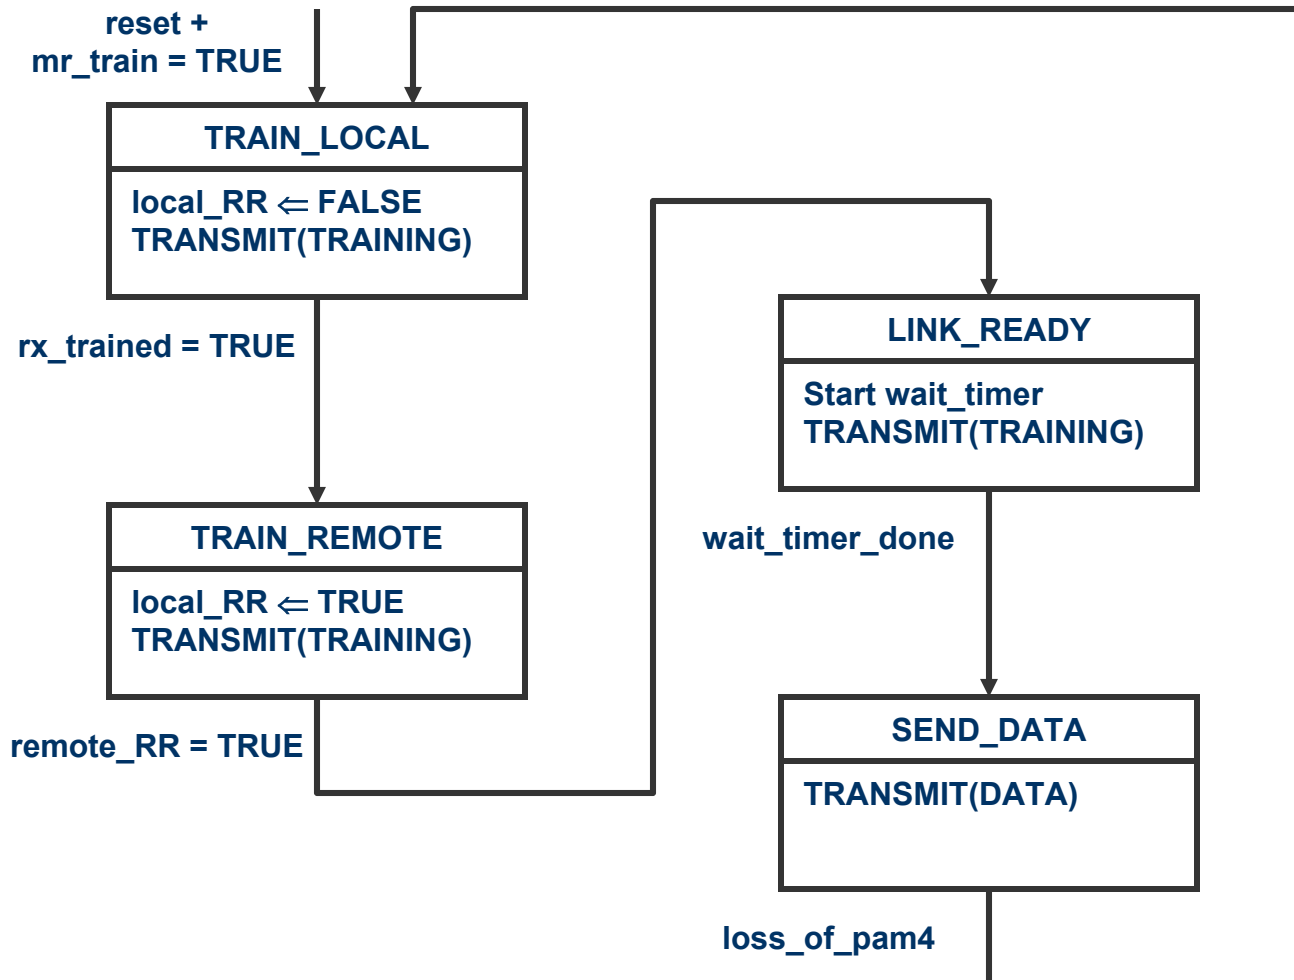
- Variables

- **reset**: Condition that is true until such time as the power supply for the device has reached its specified operating region.
- **mr_train**: Asserted by system management to initiate training.
- **local_RR**: Asserted by the link initialization protocol state machine when rx_trained is asserted. This value is transmitted as the *ReceiverReady* bit on all outgoing LIP frames.
- **remote_RR**: The value of remote_RR shall be set to FALSE upon entry into the TRAIN_LOCAL state. The value of remote_RR shall not be set to TRUE until no fewer than three consecutive LIP frames have been received with the *ReceiverReady* bit asserted.
- **rx_trained**: Asserted when the transmit and receive equalizers have been optimized and the normal data transmission may commence.
- **loss_of_pam4**: Asserted when X consecutive symbols are received without the presence of -1 or +1 symbols. This is an indication that the remote transmitter has reverted to LIP frames. The value of X shall be between 500 and 1500 PAM-4 symbols

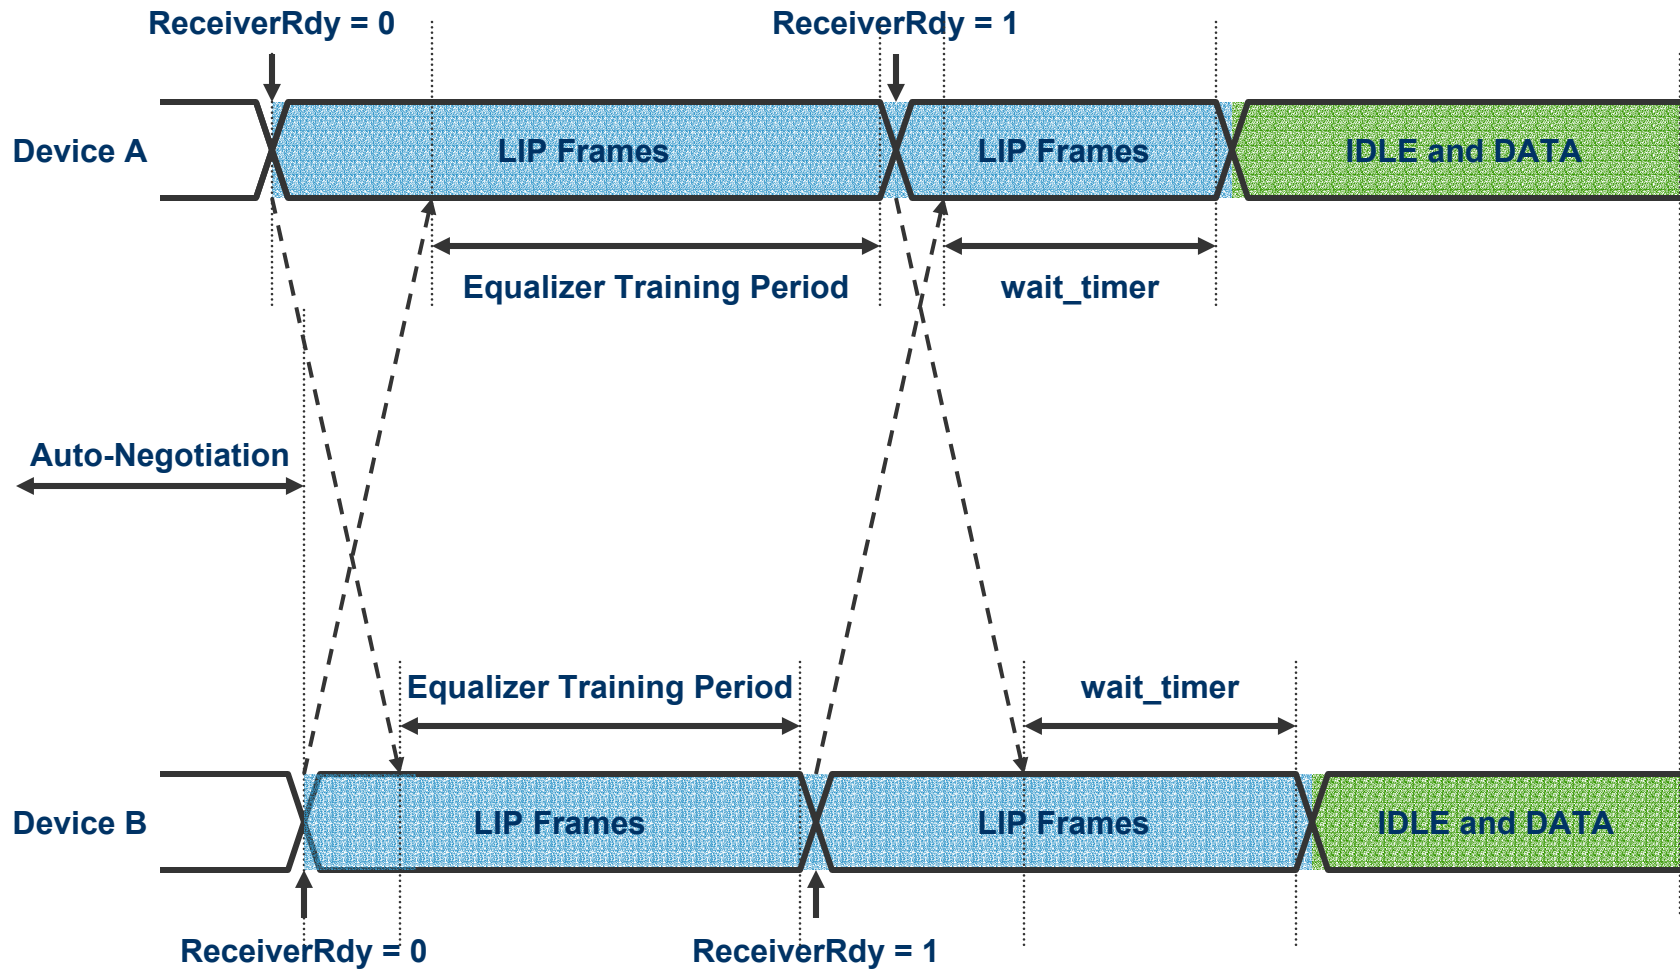
LIP State Diagram (2/3)

- Timers
 - **wait_timer**: This timer is started when the local receiver detects that the remote receiver is ready to receive PAM-4 data. The local transmitter will deliver wait_timer additional LIP frames to ensure that the remote receiver correctly detects the *ReceiverReady* state. The value of wait_timer shall be between 100 and 300 LIP frames.
- Messages
 - TRANSMIT()
 - **TRAINING**: Sequence of LIP frames. The status report and coefficient update fields are defined by receiver adaptation process.
 - **DATA**: Sequence of PAM-4 symbols as defined by the output of the PAM-4 encoding block.

LIP State Diagram (3/3)



Example LIP Timing Diagram




Note: Simulations use same worst-case channel studied earlier.


Robust Reception of LIP Frames

- LIP Frames are transmitted using PAM-2 at 5Gbaud for more reliable reception over unequalized channels.
- Robustness may be improved through the use of simple equalizer pre-sets.

No Preset




Preset for 75% de-emphasis ($c_0 = 11/14$, $c_{+1} = -3/14$)



Agenda

- Proposal Overview
- Link Simulations
- Link Initialization Protocol (LIP) Detail
- **Conclusions**


PAM-4 Silicon Complexity and Cost



Based on metrics today from devices on same 130nm process node

- Single lane of 10Gb/s PAM-4 is less than ½ die area of a typical 10Gb/s XAUI quad today.
 - Will follow similar XAUI cost declines going forward.
 - Total Cost = Chip Test + Yield + Packaging as well as Backplane Interconnect.
- 10Gb/s PAM-4 is technically feasible and demonstrated in 130nm today.
 - Extensive data for operation over 40" low-cost FR-4 backplane with two connectors.

PAM-4 Power Considerations



Based on metrics today from devices operating over 40" of low-cost FR-4 on same 130nm process node

- Single lane of 10Gb/s PAM-4 is less than $\frac{1}{2}$ the power of a typical 10Gb/s XAUI quad today.
 - Will follow additional power decline curve moving to smaller geometries.
 - This estimates includes higher voltage supplies for I/O (however, it is possible that the higher output voltage is not required for the targeted channels).

Objectives Check

- Preserve the 802.3/Ethernet frame format at the MAC Client service interface. [Yes]
- Preserve min. and max. frame size of current 802.3 Std. [Yes]
- Support existing media independent interfaces. [Yes, XGMII via the 10GBASE-R PCS]
- Support operation over a single lane across 2 connectors over copper traces on improved FR-4 for links consistent with lengths up to at least 1m. [Yes, 10Gb/s operation simulated and demonstrated]
 - Define a 1 Gb/s PHY
 - Define a 10 Gb/s PHY
- Consider auto-negotiation.
- Support BER of 10^{-12} or better. [Yes, 10Gb/s operation simulated and demonstrated to BER better than 10^{-12}]
- Meet CISPR/FCC Class A. [Automatic power control and reduction in occupied bandwidth help meet this requirement]

Conclusions

- A new PMD sublayer based on PAM-4 signaling is proposed.
- Use of transmitter pre-compensation greatly reduces receiver complexity.
- Link Initialization Protocol (LIP) maintains plug-and-play feel.
 - Simple and robust.
- Methodology proven in simulation and in measurement.
 - http://ieee802.org/3/bladesg/public/jan04/hoppin_01_0104.pdf
 - http://ieee802.org/3/bladesg/public/mar04/hoppin_01_0304.pdf
- Proposed PMD satisfies the 5 Criteria and all Task Force objectives related to the 10Gb/s serial backplane PHY.



Thank You